

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

The White House Office of Science and Technology Policy and the National Science Foundation released a Request for Information (RFI) on July 23, 2021 to inform the work of the [National AI Research Resource \(NAIRR\) Task Force](#) as they develop an NAIRR implementation road map. The RFI was published in the Federal Register and the comment period was from July 23, 2021, through October 1, 2021 ([extended](#) from an original deadline of September 1, 2021).

This document contains the 84 responses received from government, academic, and industry stakeholders. In accordance with the RFI instructions, only the first 10 pages of content were considered for each response. Please note that these responses do not represent the views or opinions of the U.S. Government or the NAIRR Task Force.

Table of Contents

- Accenture5
- The Aerospace Corporation.....12
- AI Now Institute of New York University and Data & Society Research Institute.....23
- AI Redefined, Inc.....34
- The Alexandria Archive Institute (Open Context).....40
- Amazon Web Services47
- American Civil Liberties Union (ACLU).....58
- American Psychological Association (APA).....69
- Anthropic.....74
- Argonne National Laboratory.....85
- ACM US Technology Policy Committee.....91
- Atlantic Council GeoTech Center.....99
- Michael August.....110
- BeeHero.....112
- Booz Allen Hamilton.....116
- Cadence127
- CalypsoAI Corp.....131
- Ben Freed and Howie Choset142
- Carnegie Mellon University151
- Center for Data Innovation.....160
- Center for Democracy and Technology.....171
- Cerner Corporation.....180
- Computing Community Consortium, Computing Research Association-Industry, Association for the Advancement of Artificial Intelligence.....190
- Consumer Reports.....201
- CrowdAI205
- Jared Freeman, Drew Leins, Niall Gaffney.....214
- The Data Foundation.....222
- Deloitte225
- Digital Diagnostics236
- Electronic Privacy Information Center (EPIC).....242
- Engine249

The Enterprise Neurosystem253

FABRIC Testbed264

Center for Security and Emerging Technology271

Google280

Hewlett Packard Enterprise.....291

Hyperion Research302

IBM308

IEEE Standards Association.....319

Indiana University.....323

Infiltron327

Information Technology Industry Council333

Internet2.....344

Kermit Kubitz355

Lawrence Berkeley Laboratory.....357

Lawrence Livermore National Laboratory.....368

Lawrence Berkeley National Laboratory Machine Learning Group376

Wayne Gilmore, John Goodhue, Christopher N. Hill, David Kaelli, Eric Kolaczyk, Jim Kurose, Scott Yackel.....379

Mathematica386

Medical Imaging and Resource Center, University of Chicago.....397

Microsoft409

NSF AI Institute for Artificial Intelligence and Fundamental Interactions.....420

The MITRE Corporation423

Moffitt Cancer Center.....434

NASA.....439

National Center for Atmospheric Research.....446

National Energy Technology Laboratory455

NIYAM IT, Inc.465

Representatives from the National Oceanic and Atmospheric Administration (NOAA) Artificial Intelligence Executive Committee (NAIEC) and the Center for Artificial Intelligence (NCAI)476

Noblis.....482

Northeastern University493

NVIDIA499

Open Commons Consortium at the Center for Computational Science Research, Inc.....510

Oracle America, Inc.514

Palantir Technologies, Inc.....525

Partnership on AI.....536

Maria Patterson.....545

Savash Kapoor, Mihir Kshirsagar, Arvind Narayanan550

John T. Feddema, David J. Stracuzzi, James R. Stewart.....558

SAS.....566

Abas Abdoli, Ryan N Coffee, Auralee Edelen, Michael Kagan, Daniel Ratner, Sohail Reddy, and Kazuhiro Tera579

Stanford Libraries586

Stanford University Institute for Human-Centered Artificial Intelligence (HAI).....591

U.S. Chamber of Commerce Technology Engagement Center.....602

University of Florida608

University of Illinois Chicago613

National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign.....619

NSF AI Institutes628

Thomas Yankeelov.....637

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Accenture

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



September 1, 2021

Dr. Lynne Parker

Director

National Artificial Intelligence Initiative Office, White House Office of Science and Technology Policy

Re: Implementation Plan for a National Artificial Intelligence Research Resource (Docket Number: 2021-15660)

Dear Dr. Parker,

Accenture is pleased to provide input on the implementation plan for the National Artificial Intelligence Research Resource (NAIRR).

As a leading global professional services company, Accenture provides a broad range of services and solutions in strategy and consulting, technology, interactive, and operations, that span all industries. We combine artificial intelligence (AI) with deep industry and analytics expertise to help our clients embrace these emerging, intelligent technologies confidently and responsibly.

At Accenture Labs and Accenture Federal Services, we incubate new concepts and apply the latest technologies to design and deliver breakthrough solutions for business, government, and society. In addition, Accenture's Applied Intelligence practice delivers AI applications at scale, underpinned by our Responsible AI focus on the ethical, transparent, and accountable use of AI technologies in a manner consistent with user expectations, organizational values, and societal laws and norms.

As the National Artificial Intelligence Research Resource (NAIRR) Task Force develops the implementation plan for the NAIRR, Accenture believes that the Task Force should closely coordinate with the National Science Foundation (NSF) and its National Artificial Intelligence Research Institutes. Accenture is a proud sponsor of one of these institutes – the NSF AI Institute for Adult Learning and Online Education – and believes deeply in the need for public-private collaboration on AI research.

Thank you for your work on the NAIRR, and we look forward to further participation in its development.

Sincerely,

Robert Cresanti

Government Relations Executive Director and Head of Global Government Relations Network
Accenture

Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource

1. What options should the Task Force consider for any of roadmap elements A through I above, and why?

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success

The primary goal for the NAIRR should be to serve as a champion for the sharing of data. The field of AI research would make tremendous advances if private industry and government agencies shared more data and if needed protections and systems were in place to allow for such sharing to occur. The Task Force should consider ways to incentivize increased data sharing.

Additionally, the NAIRR should endeavor to provide scalable computing resources and to create a benchmarked marketplace to advance the state of the art in AI research. When evaluating the computing resources that the NAIRR should make available to researchers, the Task Force should first seek to understand the current state of existing efforts in this space, given the fact that other government agencies like the National Laboratories already make computing resources available to researchers.

Finally, the NAIRR should seek to promote a multifaceted and multimodal approach towards AI. Recently, deep learning, or representation learning more generally, have comprised a majority of academic and industry focus. This will ultimately reach a limit. In addition to machine learning, the NAIRR should encourage researchers to focus on other areas of AI, such as:

- Decisioning and Contextual Adaptation
- Probabilistic and Symbolic Reasoning
- Causal Inferencing
- Generative Learning
- Data Efficient Learning
- Privacy Preserving Data Design and Learning
- System Engineering and Cybernetics

Some metrics that the NAIRR could use to measure success include:

- How much data is shared with the NAIRR that wasn't previously available to researchers
- How many new researchers use the advanced computing resources provided by the NAIRR and the demographic diversity of AI researchers
- How many researchers test against the NAIRR benchmarks or use their data sets

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

Some potential incentives that the Task Force could consider to encourage additional sharing of data range from support for AI-related budget requests to leadership roles in government decision-making processes. Accountability could include regular reporting requirements such as publication and promotion of statistics for leading organizations.

One potential barrier to the sharing of “high-quality” data is a lack of agreement on what exactly makes data “high quality.” When Accenture talks about “data quality,” we look at a combination of factors such as completeness, accuracy, lack of bias that negatively impacts society, relevance, and timeliness of data in relation to the insights we are trying to generate.¹ A combination of clear ownership and data lineage is how data quality is maintained and trust is built over time, and the Task Force should consider processes and frameworks around storage, management, and transfer of data that are needed to ensure “high-quality” data.

2. Which capabilities and services provided through the NAIRR should be prioritized?

The top priority for the NAIRR should be to increase access to data through the provision of high-quality, curated data sets. Of particular importance is the provision of data related to climate, healthcare, and economics – three critical fields of study that would benefit greatly from increased AI-related research.

While increasing access to advanced computing resources is a laudable goal, the National Laboratories have made considerable progress in this space. Instead of duplicating efforts, the NAIRR Task Force should work with the National Laboratories to evaluate their progress, identify gaps, and work collaboratively to complement their work.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

We believe strongly in the importance of the ethical, equitable, and fair use of technology. To reinforce these principles in the research and development of AI, organizations should adapt a two-pronged strategy:

1. Humans in the Loop: Participation in the NAIRR, should only be allowed after:
 - Appropriate terms and conditions have been attested and signed to.
 - Participants have taken mandatory training from approved providers on responsible and ethical use of the NAIRR.
 - The participating entities have a strong operational charter to mandate responsible use of data and models.

¹ “How to build a data strategy to scale AI.” Millman, Nick. May 15, 2020. <https://www.accenture.com/us-en/insights/applied-intelligence/build-data-strategy>

- Accountability from a legal and liability standpoint is established.

2. Work Products and Outputs

- The process and outputs from NAIRR research, as well as training data sets, should be auditable. Model research and development should be chronicled and available for review.
- The NAIRR should provide certain automated testing frameworks to back test the AI models against benchmarked compliance policies.
- Significant results should be reproducible.

The NAIRR can reinforce principles of ethical and responsible research and development of AI by coordinating with the National Institute of Standards and Technology (NIST). NIST has hosted numerous workshops on AI issues, such as mitigating bias in AI, and recently published a draft document titled, “A Proposal for Identifying and Managing Bias in Artificial Intelligence.”² Additionally, the Task Force should ensure that data sets have the appropriate governance in place to reduce representation bias and that all data collected from individuals should be submitted with the associated consent forms that were used to obtain informed consent and that specify the scope of the consent. For universities, this means ensuring that data sets were collected in accordance with institutional review board (IRB) criteria. For private companies that don’t have to undergo IRB, the government should ensure that data sets have been collected with documented reasonable governance with policies and processes that include informed consent.

Another approach the NAIRR could take to drive ethical behavior would be to create a robust regime analogous to the “open source” principles found in software development. The Task Force should consider requiring users who transform, add to, or modify NAIRR data to contribute such modifications back to the NAIRR. Users would have broad rights to use the data and share their use cases (models) to enhance global learning, while encouraging a community (as in open source) to contribute data updates that reduce bias and drive ethical behavior.

Finally, a way to reinforce ethical and responsible AI for deep learning (DL) models in particular is to ensure that they do not propagate negative stereotypes. This can be done by testing predictions against benchmark data or by providing a report on their testing and predictions if benchmark data does not exist. This would help address ecological fallacies, among other things.

4. *What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?*

The Department of Energy’s National Laboratories should serve as a crucial building block for the NAIRR, as they are already focused on advanced technology research and are providing

² “A Proposal for Identifying and Managing Bias in Artificial Intelligence.” Schwartz, Reva; Down, Leann; Jones, Adam; and Tabassi, Elham. June 2021. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270-draft.pdf>

researchers with access to advanced computing resources. For example, the Argonne Leadership Computing Facility (ALCF) provides researchers with access to supercomputers, visualization clusters, advanced data storage systems, and high-performance networking capabilities. According to the Department of Energy, “One important reason for establishing America’s national laboratory system immediately after World War II was to provide a home for large-scale, costly scientific facilities that universities could not afford.” Given the National Laboratories’ work in this space, the Task Force should engage extensively with the National Laboratories to determine ways to best help researchers and ensure minimal duplication of efforts.

Another effort that the NAIRR Task Force can build upon is the Coleridge Initiative’s Administrative Data Research Facility (ADRF). The ADRF is a cloud-based platform that provides secure access to government data sets that have historically been unavailable to researchers. Currently, the ADRF provides access to over 100 confidential government datasets from more than 50 different agencies. The NAIRR should consider ways to encourage government agencies to contribute more data resources to efforts like the ADRF. By leveraging an existing resource like ADRF, the Task Force can increase the speed and decrease the costs of implementing the NAIRR.

Finally, the Task Force should coordinate closely with the National Science Foundation’s AI Research Institutes. NAIRR, in concert with the NSF’s recently launched institutes, can serve as a marketplace to cross fertilize ideas across the seven research domains identified by the AI Research Institutes. This would not only accelerate research and development in AI, but it will help develop robust data sets that can be used as benchmarks for different industrial use cases.

5. *What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?*

The Task Force should consider public-private partnerships a crucial component of the NAIRR, and one model they should be considered is the aforementioned NSF AI Research Institutes. These Institutes bring together academia, industry, and government to work together and advance AI research in ways that would be impossible separately. Accenture is a major partner of NSF in these efforts, having partially funded the establishment of the NSF AI Institute for Adult Learning and Online Education. Led by the Georgia Research Alliance, this institute will seek to enhance the quality of adult online education and make education more available, affordable, and equitable. The NAIRR could follow a similar model of working with specific research institutions and industry partners to identify key areas where the NAIRR could advance AI-based fields of study.

Another model that the Task Force should examine while developing a plan for the NAIRR is the Human Genome Project which ran from 1990 to 2003. Through its sequencing of DNA, the Human Genome Project advanced a variety of scientific fields including anthropology,

biotechnology, and biofuels. The NAIRR should strive to bring together a similarly impressive group of stakeholders to advance AI research.

A final model that the Task Force should consider is the European Commission’s Destination Earth (DestinE), an effort to create a “digital twin of Earth.” This effort is intended to improve digital modelling of Earth’s physical resources, allowing researchers to simulate natural phenomena more precisely, continuously monitor the health of the planet and model the effects of climate change. DestinE will be a cloud-based modelling and simulation platform, where users will be able to access data, advanced computing infrastructure, software, AI applications, and analytics. Given the similarities between DestinE and NAIRR, the Task Force would be well-served by closely examining DestinE’s design and implementation.

6. *Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?*

One of the major challenges that the NAIRR will face is the question of liability. If NAIRR succeeds in encouraging federal agencies, private companies, non-profit organizations, and others to provide data to this effort, NAIRR will need to provide clear guidance on the potential liability concerns surrounding shared data. Without clarity and confidence, non-government entities will be reticent to participate, potentially creating a significant limitation on the potential for NAIRR’s success.

The Task Force will need to consider how best to prevent the unethical use and sharing of data collected from individuals both for its own altruistic merits and as an incentive for responsible entities to participate. Relatedly, the NAIRR will have to negotiate difficult questions around the extent to which consumers can demand recourse.

A final limitation that NAIRR will face in democratizing access to AI research and development is the overall lack of diversity within the AI field. From race and gender to academic training and vocational background, AI, and the technology field more generally, has a well-documented diversity problem. Organizations like Girls Who Code and AI4ALL have been founded to help close various diversity gaps and increase representation in the technology field. In developing a plan for the NAIRR, the Task Force should ensure that diversity and inclusion are at the center of their efforts.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

The Aerospace Corporation

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Attachment A: Aerospace Response and input to RFI questions

1. What options should the Task Force consider for any of roadmap elements A through I above (from the RFI), and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success; The goals and metrics for this endeavor should state measurable, to the fullest extent possible, strategic end states that provide clarity and substance to outcomes. Characteristics of successful metrics are described below:

- Metrics are key to
 - *Evaluating decision points*
 - *Making a defensible argument*
- There are two different families of metrics
 - *Model Metrics*
 - Evaluating a machine model is a key aspect of determining whether it is useful in answering key business and operational questions. In this regard, this regard there is a fundamental tradeoff in performance between:
 - *Evaluation—how accurate a model is making inference (based on training and validation data)*
 - *Generalizability—how generalizable or scalable a model is (based on hold out data or ground truth)*
 - *Benchmarking—how well a model compares to other machine learning models*
 - Evaluation is intimately tied to the type of machine learning problem you are trying to solve
 - *Enterprise Metrics*
 - In order to properly track usage, derive effective insight on product usage, and deliver enterprise value, it is imperative to develop appropriate business-level metrics
 - There are two flavors of enterprise-level metrics:
 - *User Metrics: track service usage and provisioning behaviors*

- *Enterprise Metrics: track adoption across the national enterprise*

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and

ii. A governance structure for the Research Resource, including oversight and decision-making authorities.

This plan should state clear roles and missions of Departments and Agencies including appropriations, authorities, and responsibilities up front to ensure maximum efficiency.

- *Emphasis of organization should focus on fostering innovation through sponsorship of cross-cutting capabilities*
 - Through coordination of funding
 - Making available large repositories for model training
 - Encourage sharing of information, data, code, best practices through online digital collaboration platforms, conferences, and representation on standards bodies
- *Aperture opens to include AI/ML within the Federal enterprise, open source, commercial, and academia*
- *Imperative to partner across USG and look for on-ramps*

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources.

Governance and Oversight structure and process should be specified in writing to establish strategic direction, make programmatic decisions, and manage the allocation of resources. Conceptual models and processes for the National Artificial Intelligence Research Resource (NAIRR) are depicted in the two figures below.

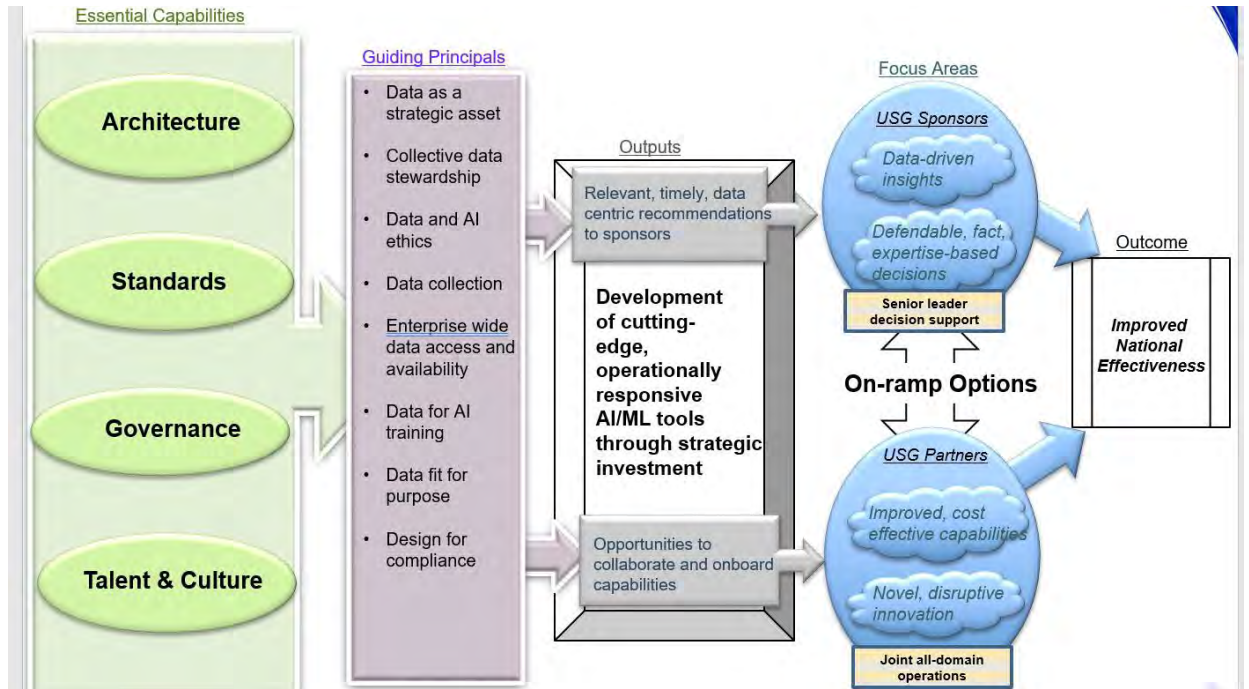


Figure 1: Holistic flow of inputs and decisions




Figure 2: Partnership Ecosystem

NAIRR will:

- Support Academia in development of long-term AI/ML capabilities as well as scan the horizon for potential technology disruption
- Partner with NSF to identify, evaluate, and support funding opportunities for AI/ML research in the national interest
- Work with Department of Commerce to develop incentivization schema to foster US AI/ML business growth
- Collaborate with standards bodies to create systematic and repeatable metrics to evaluate AI/ML performance
- Advise and work with department of state to “hit the sweet spot” between specifying ITAR restrictions necessary to safeguard technologies critical to the national interest and cross-national collaboration
- Ensure that national security requirements are met through investment and premonition of AI/ML investments and technology across federal, commercial, and academic realms
- Partner with Venture capital to identify and share investment options in develop of organic AI/ML startups

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advance computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure. Considerations of this characteristic would consist of, but not be limited to:

- Different combinations of data produce different analytics products
- The second most labor-intensive aspect of AI/ML development is curating data
- There are several nuances involved in data cleaning, sanitation, formatting etc.
- Having pre curated data is a boon for developers
- Linking these datasets (curated) to specific AI/ML problem areas and use cases would greatly facilitate knowledge discovery and make it efficient for developers to find what they are looking for
- Few, if any, analytics vendors keep their own infrastructure
 - *Cost is too high (especially for start-ups trying to hit VC backed revenue goals)*
 - *Not scalable to the really hard problems*
- Having a Graphics Processing Unit (GPU) enabled cloud infrastructure with persistence storage provided by NAIRR would be extremely helpful to developers as it would:
 - Give the ability to train models of appreciable complexity
 - Allow developers to devote funds to other activities (such as hiring researchers or funding other projects)
- In terms of infrastructure, providing support for containerization and orchestration is a key enabler of AI/ML DevOps

- It also ensures stability and consistency of deployment
- Providing a platform to share code, allow comments and exchange between developers, and promotion of a leaderboard across the variety of different AI/ML problems would greatly facilitate collaboration and also provide a resource from which developers could learn and apply to new problems
- Proper security access controls ensure that the NAIRR resource cannot be manipulated or otherwise repurposed to support the aims of malevolent actors

E: An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the NAIRR.

- There are numerous case studies/publications citing why AI products fail
- 85% fail due a handful of reasons (as shown below):

AI/ML Adoption Barriers

Which of the following has been the greatest barrier to adopting AI and machine learning in your organisation?





Figure 3: Barriers to AI Adoption

AI/ML Challenges



Which of the following are the top challenges when developing your AI application?



Source: Gartner 2020 AI Review

Figure 4: Top challenges to developing AI applications

Other Risks/Challenges:

- Data Science initial models don't scale or are too experimental to be used by internal or external customers.
- Data science/ML models while brilliant and innovative don't meet the business requirements or are too fragile to respond to change in the supporting data.
- AI initiatives are driven by the company's internal IT organizations and inherit "waterfall" challenges.
- Companies don't have the patience for the time it takes to deliver on AI/ML projects.
- A lack of a "Product" approach to AI/ML projects is core to project failure and increased risk.

Mitigation to these risks:

- Anticipating and planning against such risks only serve to strengthen the national interest

- It also helps provide insights as to when it's appropriate to either onboard or offboard AI/ML investments or otherwise course correct policy
- Once risks are identified, they must be addressed, though, some level of risk is necessary to innovate
- The question becomes how to balance risk versus reward
- Such guidance is key to building a roadmap

F: An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls.

Cybersecurity, including access controls are an imperative and should be addressed up front. Based on extensive DoD Cybersecurity score carding, a significant percentage of incidents are due to Two Factor Authentication, Phishing, and Insider-enabled events.

Basic Cyber Hygiene to address these attack vectors is vital to maintaining the security and resilience of this initiative, which, if there is an attack, will receive significant attention. National resources in the Intelligence Community, the Department of Defense, and Federal Law Enforcement should be brought to bear for overall strategic oversight and defense against the small, but significant portion of threat events enabled by foreign nation state capabilities.

G: An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research.

Privacy and civil rights and civil liberties are foundational to this effort and should receive primary leadership and staff attention to ensure they are holistically included in the entire NAIRR effort.

H: A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; A clear, multi-year research and development investment profile, itemized by major program elements and further apportioned by Department and Agency is cross cutting and essential for effective governance.

- Identification and pursuit of funding opportunities and partnerships is critical to the NAIRR sustainment
- Because of funding limitations, the sheer complexity of AI/ML challenges, and the pace in which the landscape is changing no organization can provide a national AI/ML advantage if acting in isolation
- A partnership is an opportunity for NAIRR to align with another USG agency, company, or investment firm to support the development or deployment (or sometimes both) of a critical AI/ML capability

- Partnering improves both potential technology reach and joint operational effectiveness, but it also introduces a few complexities for development and deployment
- There are many different types of funding & partnering options, each with associated tradeoffs
- This is not a one size fits all.
- Different options must be exercised based on circumstance
- The key is conducting the requisite due diligence to find which option is appropriate (to mitigate risk and ensure an optimal outcome) and then structure a portfolio of these options in order to promote national AI/ML interests

Partnering Options

Option	Benefit	Risk
NAIRR spearheads development and owns/leverages transition	Easier to manage. Fewer interfaces and decision gates required for development and deployment.	Data availability and funding risk may impair efficacy. Harder to guarantee joint all-domain operational effectiveness. Transition predicated on fewer end-users.
NAIRR spearheads development and transitions to partner	Easier to manage. Fewer interfaces and decision gates required for development. Moderate potential for joint all-domain effectiveness.	Data availability. Funding and transition risk dependent on partner disposition.
Partner spearheads development and NAIRR owns/leverages transition	Potential capability add and/or integration to tech stack. Fewer interfaces and decision gates required for deployment. Moderate potential for joint all-domain effectiveness.	Harder to manage. Viability of technology solution dependent on partner disposition.
NAIRR collaborates with collaborates in development and partner owns/leverages transition	Potential capability add and/or integration to tech stack. High potential for joint all-domain effectiveness.	Harder to manage. More interfaces and decision gates required for development. Viability of technology solution dependent on partner disposition.
NAIRR collaborates with partners in development and partners own/leverages transition	High potential for joint all-domain effectiveness.	Harder to manage. More interfaces and decision gates required for development. Viability of technology solution dependent on partner disposition. Funding and transition risk dependent on partner disposition.
NAIRR collaborates with partners in development and consortium owns/leverages transition	Potential capability add and/or integration to tech stack. Highest potential for joint all-domain effectiveness.	Hardest to manage. More interfaces and decision gates required for development and deployment. Viability of technology solution dependent on partner disposition. Funding and transition risk dependent on partner disposition.

Due diligence is the key to mitigating risk and ensuring optimal outcomes

Figure 5: NAIRR Partnering Opportunities

I: Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

- Agency Roles and Responsibilities addressed in (A).
- What’s important for (I) is that NAIRR must “own” the due diligence process, determine which partnering/investment options are appropriate, and then structure a portfolio of options to serve the national AI/ML interest.

- Aerospace suggests the following due diligence process.

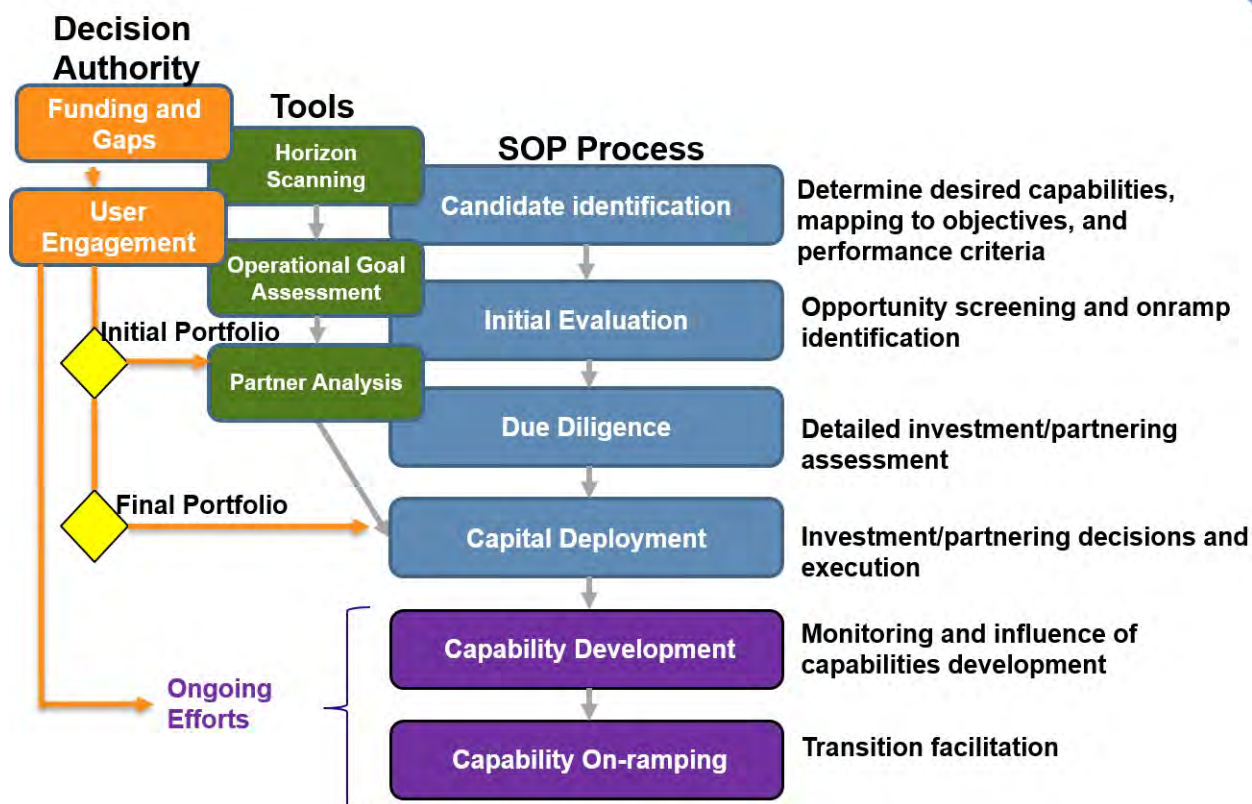


Figure 6: NAIRR Due Diligence Process

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Aerospace recommends prioritizing in this order: B, A, C, D, with the other items being specific annexes and sub-initiatives to C.

Furthermore, security access controls are considered a prerequisite for any front facing platform deployment. Following this, in order of precedence:

- Data Repository
- Data Curation
- Infrastructure
- Code repository
- Developer Engagement Platform
- Leaderboard
- Onboarding of additional learning resources

Critical to this is the development and integration of enterprise metrics in order to understand what elements of the NAIRR platform are most used, under-used, etc. A/B testing and user outreach should be integrated into DevOPs so that features most useful to end-users are offered.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Ethical implications of AI include liability and law and are necessary conditions to integrate moral, societal, and legal values with technological developments in AI. Responsible AI is a priority and this theme and responsibility should be a foundational part of Item C.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Broad existing initiatives exist in the following areas that should be leveraged to the maximum extent possible:

- Quantum and High-Speed Computing
- Cybersecurity – Specifically the Comprehensive National Cybersecurity Initiative (CNCI) from 2007 – 2014. 12 Initiatives plus Enablers – strong model for this AI initiative
- All recent NIST Special Publications on AI, SP 1270 Identifying and Managing Bias in Artificial Intelligence
- Advanced Autonomy
- Private Sector Big Data Analytics
- Advanced Manufacturing
- Trusted Micro Electronics
- Innovation (for example, the Department of Defense’s Developmental Innovation Unit (DIU) in Mountain View, CA)

Several building blocks exist that can be leveraged. In terms of company selection for partnering, both Crunchbase and Pitchbook maintain voluminous stores of information that can be accessed either online or via API call. The Microsoft Academic Graph has data on over 500 million academic publications that can be mined for the purpose of technology identification, horizon scanning, or research collaboration.

IEEE and ACM have active AI/ML panels that could be tapped into to provide situational awareness. NSF could also be tapped into as well. Many academic AI/ML institutions almost have incubators/angel investment arms that could help provide targeted engagement for investment sharing. In-Q-Tel and DIUX could likewise provide investment intelligence from the standpoint of alignment with national security AI/ML requirements.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

AI Now Institute of New York University and Data & Society Research Institute

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

October 1, 2021

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource

Dear Members of the National AI Research Resource Task Force,

The AI Now Institute of New York University and Data & Society Research Institute are pleased to submit a response to the Request for Information (RFI) published by the Office of Science and Technology Policy and the National Science Foundation (NSF) to inform the work of the National Artificial Intelligence Research Resource (NAIRR) Task Force.

Our organizations are interdisciplinary research institutes working to help ensure that artificial intelligence (AI) systems are accountable to the communities and contexts in which they are applied, and to produce empirical research that challenges the power asymmetries created and amplified by technology in society. We have worked extensively within academic institutions, civil society and advocacy communities, and in solidarity with marginalized communities and workers directly affected by algorithmic harms.

We are pleased to see the government's commitment to supporting research on AI, and for the opportunity to contribute to the development of the Task Force's implementation blueprint. We are filing this comment to express our concern about the stated aims for the NAIRR and to recommend alternative policy strategies for supporting research into AI and expanding access to data, resources, educational opportunities, and meaningful mechanisms that can ensure democratic oversight of AI and related technologies. We argue that these changes are needed to enable high-quality, interdisciplinary work that goes well beyond a narrow technocratic frame.

Throughout our comment, our recommendations encourage the Task Force to fundamentally reconsider whether the investment in shared computing and data infrastructure is consistent with both the broader aims of the project to “democratize” AI research and the Biden Administration’s explicit commitment to challenging the concentrated power of the tech industry.

In encouraging the Task Force to reconsider, we emphasize that the only plausible short- to mid-term scenario is that the infrastructure required for the NAIRR would be licensed from the same large tech companies responsible for concentrating tech power. This means that while the NAIRR proposal claims to “democratize” access to these resources as a way of contending with the concentrated power of the companies that control AI infrastructures, it would in reality almost certainly work to expand and entrench the power and control these companies have over these

infrastructures.

The NAIRR proposal will make it much harder to check the power of these companies through regulation and public pressure. Furthermore, this proposal assumes a narrow understanding of the AI research field and what kinds of resources are necessary to do the work that's most needed to not only deepen our understanding of AI, but to prevent and mitigate harms that AI and the industry concentration behind it are already causing in our society.

These are challenging issues for which there are no current best practices. Our comment urges the NAIRR Task Force to recommend that Congress consider alternative spending priorities and deeper engagement on the harms unleashed by uncritical investment in AI.

In this comment, we argue:

1. [The NAIRR will entrench, rather than challenge, corporate control over the AI field, contrary to the Biden Administration's bold stance against the power of large tech companies in society.](#)
2. [The NAIRR Task Force must reckon with mounting evidence of the harmful impacts of large-scale AI systems, including discriminatory consequences for marginalized groups and long-term climate impact.](#)
3. [The NAIRR Task Force must expand their understanding of what disciplines constitute "AI research" and redirect NSF resources and programs toward constructing mechanisms for meaningful democratic control of AI and related technologies.](#)
4. [The NAIRR and related proposals raise serious ethical and data privacy challenges, particularly with the use of government data. Given the lack of demonstrated best practices and global policy precedents for data privacy, the NAIRR Task Force should recommend pausing the NAIRR until these challenges are resolved.](#)
5. [The NAIRR, as presently conceived, bolsters misleading and dangerous "tech cold war" narratives, which reflect the self-interest of Big Tech and the defense contracting industry, without being backed by robust evidence.](#)

Background and context: Understanding the "AI boom" (*Question 1, Items A and H*)

To understand the complex interplay between the issues the NAIRR proposal raises, it's helpful to recall that the current turn to AI is primarily a product of significantly concentrated corporate resources—namely vast computation, massive data, and the capital required to attract and retain

scarce AI talent.¹ **The so-called “advances” in AI that have been celebrated since the early 2010s were not due to breakthroughs in AI research and innovation. They were predicated on newly available access to powerful computation and to massive amounts of web data.** Then, as now, these are resources that a handful of powerful tech companies have in large supply thanks to ad tech-driven surveillance business models, and that few others can avail themselves of without going through these companies.²

The past decade’s Big Tech-led turn to AI profoundly shaped academic computer science disciplines as well. It served to redirect computer science research toward AI-related questions and approaches favored by these companies. In particular, the influx of money and attention produced a turn toward resource-intensive research and development. Work that could avail itself of expensive and scarce industry computing and data was heralded as “cutting edge.” This created an uneasy and conflicted environment for university AI research, in which the dependence on large tech company funding, infrastructure, and data was recognized by practitioners, but not often openly acknowledged.

The NAIRR proposal’s acknowledgement that all of America’s AI researchers cannot “fully participate in”³ this field is a step forward, allowing us to recognize that concentrated corporate power and resources are constitutive of the current wave of AI development. It also presents us with a set of thorny questions, at the center of which is the question of how we reduce the power and control of the handful of companies currently dominating AI and AI research, and how we ensure that determinations about whether—if at all—AI is developed and deployed are subject to more democratic deliberation. While the NAIRR proposal doesn’t currently address these issues, and indeed threatens to exacerbate them, our hope is that the Task Force can redirect attention and resources to this pressing task.

I. The NAIRR will entrench, rather than counter, corporate control over the AI field, contrary to the Biden Administration’s bold stance against the power of large tech companies in society. (Question 1, Topics H and D; Questions 5 and 6)

Although the NAIRR is currently envisioned as “a shared computing and data infrastructure that would provide AI researchers and students across scientific fields with access to a holistic advanced computing ecosystem,” it remains unclear how the NAIRR can deliver on the mandate to build an implementation plan in a way that does not further entrench corporate influence and control over the

¹ While the field of Artificial Intelligence is ostensibly oriented around making machines intelligent, in practice, most AI systems rely on big data - the collection and processing of massive datasets, identifying patterns and probabilities within them and codifying them into a predictive mathematical model. See Meredith Broussard (2018) *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge: MIT Press.

² Meredith Whittaker, “The Steep Cost of Capture,” *ACM Interactions*, Vol. XXVIII.6 Nov-Dec 2021. Forthcoming.

³<https://www.federalregister.gov/documents/2021/07/23/2021-15660/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>

AI research field.⁴ What is being proposed is an extension of industry-dependent resources, not the construction of resources that would challenge or reduce the centralized power of the large tech players.

From this perspective, it makes sense that the National Security Commission on Artificial Intelligence (NSCAI), a very conflicted body that is helmed by former Google CEO Eric Schmidt, and populated by executives from Amazon, Microsoft, Oracle, and other large tech companies, would prioritize support for such an endeavor.

There is no scenario in the short or mid-term future where large scale computational resources adequate to the task of expanding access to bigger-is-better AI research resources could be created and maintained by institutions meaningfully separate from the large tech platform companies. These companies provide more than raw computing power: the computational environments they own and license provide the tools and research environments that define how AI research gets done. There is no plausible path forward in which such a resource would not be dependent on existing tech industry platforms, tools and resources.

While this isn't stated explicitly, it is tacitly acknowledged in the repeated calls to constitute the NAIRR via "public-private partnerships," a phrase threaded throughout the Stanford Institute for Human-Centered Artificial Intelligence (HAI) and NSCAI's NAIRR advocacy.⁵ NSCAI's recommendations acknowledge that "this infrastructure would leverage public-private partnerships and build on existing government efforts, avoiding high start-up costs of a government-run data center."⁶ Speaking with *Science*, HAI Director John Etchemendy makes the implicit explicit: "The commercial cloud providers are doing the innovation, and they invest massive amounts of money to keep it up-to-date," he says. "It would be a huge mistake to build a facility like a supercomputer center because it would be obsolete within a few years."⁷

These statements make clear to us that the NAIRR proposal is a tacit endorsement of a massive investment in large tech companies in service of expanding their proprietary infrastructure. This arrives at a time when the concentrated power and influence these companies exert is increasingly under scrutiny, including by the Biden Administration itself, which has taken a clear stance that a small number of dominant platforms are using their power to extract monopoly profits.⁸ An effort that aims to "democratize" AI research by investing money in companies that

⁴<https://www.federalregister.gov/documents/2021/07/23/2021-15660/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>

⁵ <https://hai.stanford.edu/national-research-cloud-joint-letter>

⁶ <https://www.nscai.gov/2020/04/01/nscai-submits-first-quarter-recommendations-to-congress-2/>

⁷ <https://www.science.org/news/2021/01/us-law-sets-stage-boost-artificial-intelligence-research>

⁸ <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/07/09/executive-order-on-promoting-competition-in-the-american-economy/>

dominate their market will only further entrench these firms' power and reach. This will make it harder to check the power of these companies through regulation and public pressure.

We therefore advise that the NAIRR's forthcoming roadmap and implementation plans recommend against pursuing shared research infrastructure, and instead explore alternative ideas for expanding research into AI, increasing government investment in critical work on AI, and meaningfully democratizing decision making on the development and deployment of AI and related technologies.

II. The NAIRR must reckon with mounting evidence of the harmful impacts of large-scale AI systems, including discriminatory consequences for marginalized groups and long-term climate impact of this scale of computing. (*Question 1, Topics E and G; Question 3*)

The NAIRR Task Force must engage more deeply with the body of critical research, press coverage and investigative reporting, and public discussion that has presented significant evidence of AI's harms, and raised fundamental questions about the ability of AI systems to operate safely and transparently in sensitive social domains. Painting a narrow definition of technical AI research and development as imperative to national success and wellbeing ignores this robust public debate and accompanying evidence, and forecloses discussion about whether society wants these systems in the first place.

While proponents of the NAIRR and the expanding use of AI more generally often point to the potential for this technology to stimulate economic growth, many people—particularly marginalized communities—are already subject to the worst excesses, mistakes, and harms perpetuated by the oppressive and extractive use of powerful algorithmic technology.⁹ Focusing on propelling increasingly unequal economic growth while these harms continue relatively unchecked, raises critical questions about whether certain communities are considered expendable and what kinds of harms are allowable in our society in the name of economic growth.

In a joint letter addressed to OSTP in July 2021, a coalition of civil rights and technology organizations called on the department to center civil rights concerns in AI and technology policy, emphasizing the need to fully incorporate the Biden Administration's Executive Order on Racial Equity into its AI policy priorities. This letter drew on years of evidence about the harms that AI is already causing, including perpetuating housing, financial services, and hiring discrimination.¹⁰ This

⁹ <https://nyupress.org/9781479837243/algorithms-of-oppression/>;
[https://www.wiley.com/en-us/Race+After+Technology:+Abolitionist+Tools+for+the+New+Jim+Code-p-9781509526437#:~:text=Presenting%20the%20concept%20of%20the,ultimately%20doing%20quite%20the%20opposite.](https://www.wiley.com/en-us/Race+After+Technology:+Abolitionist+Tools+for+the+New+Jim+Code-p-9781509526437#:~:text=Presenting%20the%20concept%20of%20the,ultimately%20doing%20quite%20the%20opposite.;);
<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

¹⁰ <https://www.aclu.org/news/privacy-technology/how-artificial-intelligence-can-deepen-racial-and-economic-inequities>

letter also calls attention to the lack of civil rights representation on the NAIRR Task Force, despite the enormous implications of the Task Force’s work on civil rights and civil liberties.

Furthermore, many subcategories of AI, especially those that emphasize large-scale data and computing (often called Large Language Models or LLMs), are uniquely prone to perpetuating social harms and entrenching biases.¹¹ These large-scale models, which are trained on troves of internet data from sources like Reddit, exhibit persistent discriminatory outputs.¹² Furthermore, large-scale AI models are built on mass surveillance, which disproportionately impacts marginalized communities,¹³ without implementing meaningful mechanisms for accountability or consent of the public.¹⁴ Their carbon cost is substantial: the amount of processing required to train AI models is both financially and environmentally resource-intensive, and these costs are only likely to expand given industry standards that tie performance metrics to the size of the dataset used to train the model.¹⁵ As a whole, the tech industry is responsible for a global carbon footprint comparable to the aviation industry, and data centers make up 45% of this footprint.¹⁶

Rigorous empirical work to understand ways AI systems and the concentrated power of the companies behind them exacerbate societal harms is happening in a wide range of disciplines and contexts. While this work has helped inform the public and raise urgent questions, it is rarely included in definitions of fundamental AI research and development (R&D), despite how critical it is to our understanding of the feasibility of ideas like the NAIRR and ultimately the tradeoffs of greater adoption of AI in all sectors of our economy. More meaningful public deliberation around these tradeoffs is important for the NAIRR Task Force to pursue, as is additional funding and resources for AI researchers and communities exploring and exposing these harms.

III. The NAIRR Task Force must expand their understanding of what disciplines constitute “AI research” and redirect NSF resources and programs toward constructing mechanisms for meaningful democratic control of AI and related technologies.

(Question 1, Topic I; Questions 3, 4 and 6)

The proposal offers a narrow interpretation of “AI research” by emphasizing infrastructure required for AI research that is reliant on large datasets and massive computational power. In this vision¹⁷, all

¹¹ <https://dl.acm.org/doi/10.1145/3442188.3445922>

¹² <https://arxiv.org/pdf/2101.05783v1.pdf>

¹³ <https://www.law.georgetown.edu/privacy-technology-center/events/color-of-surveillance-2019/>

¹⁴ <https://dl.acm.org/doi/10.1145/3442188.3445922>

¹⁵ <https://arxiv.org/pdf/1907.10597.pdf>

¹⁶ <https://medium.com/@AINowInstitute/ai-and-climate-change-how-theyre-connected-and-what-we-can-do-about-it-6aa8d0f5b32c>

¹⁷ Proponents of the development of the National Research Cloud (NRC) such as Stanford’s HAI also put forward a similarly narrow interpretation of “AI research”.

AI research is “technical,” and research at the cutting edge is that which is most dependent on large datasets and massive computational power.

Although the NAIRR’s vision includes “AI researchers and students across scientific fields,” it’s important to be explicit about the wide range of people, organizations, and disciplines that also contribute to AI R&D. Major contributions to this field are coming from scholars in disciplines like law,¹⁸ anthropology,¹⁹ and history,²⁰ and also from community organizations²¹ gathering qualitative evidence and elevating the lived experience of the people who are surveilled, assessed, and otherwise subject to AI’s determinations and predictions. These disciplines and approaches are most capable of analyzing and addressing the structural issues with AI and related technologies.

Using this more accurate and expansive understanding of the academic disciplines and organizations leading research on AI, the Task Force must consider a wider range of needs than those required by scholars in fields such as machine learning and neuroscience. The Task Force must also deepen its understanding of what it means to “democratize” AI by centering the perspectives and needs of people and organizations whose work concerns AI but does not rely on computing or quantitative data infrastructure.

Despite the NAIRR proposal’s efforts to “democratize access to the cyberinfrastructure that fuels AI research and development,” the infrastructure that would constitute a NAIRR system would almost certainly to be licensed from the same companies responsible for concentrating AI power.

This raises a critically important question: where else could hundreds of millions of dollars be better spent, and how could such funding help us meaningfully ensure democratic control and deliberation over AI and the tech companies responsible?

The NSF has already laid important groundwork to invest in the creation of new research institutes, investing a combined \$220 million in 40 states and in the District of Columbia.²² The NSF can continue to expand this program by substantively funding under-resourced research domains, taking leadership from those most harmed by inequitable uses of AI, and investing in meaningful public control over these powerful technologies. **Rather than routing millions of dollars in taxpayer money to private industry under the guise of public-private partnership, this funding could be**

¹⁸ <https://datasociety.net/library/poverty-lawgorithms/>; <https://ainowinstitute.org/litigatingalgorithms-2019-us.pdf>; <https://www.nyulawreview.org/online-features/dirty-data-bad-predictions-how-civil-rights-violations-impact-police-data-predictive-policing-systems-and-justice/>

¹⁹ <https://datasociety.net/library/repairing-innovation/>; <https://global.oup.com/academic/product/predict-and-surveil-9780190684099?cc=us&lang=en&>

²⁰ <https://datasociety.net/library/house-arrest/>; Joy Lisi Rankin (2018) *A People's History of Computing in the United States*, Harvard University Press; Mar Hicks (2017)

²¹ <https://movementalliance.org/about/>, <https://d4bl.org/>, <https://detroitcommunitytech.org/>, <https://www.drivers-united.org/>; <https://mijente.net/2018/10/whos-behind-ice-the-tech-companies-fueling-deportations/>; <http://www.awoodcenter.org/>; <https://athenaforall.org/>; <https://www.fightforthefuture.org/>

²² https://www.nsf.gov/news/news_summ.jsp?cntn_id=303176

used to create and fund scholarships for underrepresented students and sociotechnical research disciplines, to invest in fellowships that place socio-technical scholars in federal agencies, and to create public engagement and consultation fora that allow communities most harmed by these systems to have a say in their design and deployment. The NSF must also work to preserve the independence of research undertaken in these institutes by eliminating the model in which these institutes are jointly funded by companies.

IV. The NAIRR and related proposals raise serious ethical and data privacy challenges, particularly with the use of government data. Given the lack of demonstrated best practices and global policy precedent, the NAIRR should pause until these issues are resolved and prioritize these issues for further research. (*Question 1, Topics E and G; Question 3*)

The creation of a National Research Cloud presents significant challenges to ensuring the ethical use of data. As over a decade of work on developing ethical guidelines for the use of open data indicates, these are hard problems. But given the likelihood that the NAIRR will provide access to government data sets, it is critical that there are stringent mechanisms in place to ensure that uses of this data are accountable and ethical, *before* the NAIRR is implemented. NAIRR documentation thus far does not indicate sufficient resources for or attention to ethical and privacy considerations.

There are several potential issues that must be addressed as a baseline by such a proposal if it proceeds with a focus on technical infrastructure. First, there should be some process for evaluating appropriate uses of the NAIRR resources and their potential downstream impact. Already, leading AI research conferences like NeurIPS require researchers to, at a minimum, address the social impact of their work prior to publishing their research findings.²³ We would suggest such a statement be treated as the floor rather than the ceiling for ethics evaluation.

There would also need to be an appropriate method for evaluating proposals. It is important that the actors involved in evaluation have expertise, independence, authority, and most importantly, adequate resources to conduct review of proposals for access, as well as monitor their use once granted. Though the NAIRR provides a development and research environment, any ethical guidelines implemented for access would do little to inhibit downstream harms in deployment, and AI models trained using NAIRR resources could be redeployed and commercialized in applications that depart from what was originally intended. This is a problem that is exacerbated by corporate secrecy, which allows companies to re-deploy academic models without transparency or oversight.

The NAIRR proposal also emphasizes that beneficiaries of these grants could receive access to “high quality government datasets” alongside access to compute. Increasing access to government data,

²³ <https://neuripsconf.medium.com/getting-started-with-neurips-2020-e350f9b39c28>

including anonymized or non-personal datasets, raises persistent data privacy and security concerns, alongside concerns about potential downstream racial and gender bias and inaccuracy.²⁴ While there is active and evolving research around differential privacy and other means to ensure privacy-preserving data disclosures,²⁵ there is a significant gap between theory and practice. Globally, too, this is an area of active policy making²⁶ as evidenced by the European Union’s Data Governance Act (DGA) of 2020 that puts in place foundational norms around tiered access to government data.²⁷ While the DGA stipulates that government agencies can share datasets with private actors “within a secure processing environment” that remains in the control of the public sector, the specifics of how this proposal would be operationalized is yet to be determined, and there are no working precedents to draw from.²⁸

These are all challenging issues for which there are no current best practices nor global policy precedent, and that require deep study and consideration. On this basis, the NAIRR should not proceed as it is currently envisioned. If it does proceed, it must have a clear and evidence-based demonstration that these issues can be addressed.

V. The NAIRR, as presently conceived, bolsters misleading and dangerous “tech cold war” narratives, which reflect the self-interest of Big Tech and the defense contracting industry, without being backed by robust evidence. (Question 1, Topics A and E; Question 6)

Proponents of the NAIRR have lauded the project as being critical to ‘winning’ a global technological race.²⁹ Evoking calls to nationalism,³⁰ these narratives legitimize continued investment in building out the most computationally intensive forms of AI research as a key element of competing with foreign adversaries, particularly China. **Tech company CEOs and former national security officials have been some of the most vocal in endorsing this self-serving narrative, which troublingly echoes cold war framings that served in the past³¹ to accelerate government investment in weapons-related computing.**³² As we have argued, however, this “tech cold war” narrative propels a race to the bottom, where any calls for democratic restraint and regulation of the largest American AI companies is seen as hurting national interest.³³ Indeed, we increasingly see

²⁴ <https://www.cs.princeton.edu/~arvindn/publications/precautionary.pdf>

²⁵ https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3329330

²⁶ https://law.yale.edu/sites/default/files/area/center/china/document/shifting_narratives.pdf

²⁷ <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020PC0767>

²⁸ <https://digital-strategy.ec.europa.eu/en/policies/data-governance-act>

²⁹ <https://www.cnas.org/publications/commentary/a-national-cloud-for-all>

³⁰ <https://foreignpolicy.com/2020/08/27/china-tech-facebook-google/>

³¹ <https://online.ucpress.edu/phr/article-abstract/88/4/619/80374/Artificial-Intelligence-and-Japan-s-Fifth>

³² Seymour Melman, *Pentagon Capitalism: The Political Economy of War*, New York, McGraw-Hill, 1970; Paul Edwards, “The Closed World: Computers and the Politics of Discourse in Cold War America” The MIT Press, 1997

³³ <https://medium.com/@AINowInstitute/china-in-global-tech-discourse-2524017ca856>

China-centered arguments against tech accountability and antitrust³⁴ that frame these regulatory interventions as a barrier to national progress in this so-called “AI arms race.”³⁵

The NAIRR, as it is presently conceived, bolsters this misleading and dangerous narrative by legitimizing the consolidation of the AI industry as a foregone conclusion. It not only affirms continued investment in building out computationally intensive AI as a given for national progress, but also further entrenches dependencies on existing AI infrastructure companies like Amazon, Microsoft, Google, and IBM, who unsurprisingly have been among the most vocal proponents. A senior Amazon Web Services executive for example applauded the NAIRR proposal while emphasizing that the AWS Cloud would “increase efficiency while providing unrivaled scalability” to the project. This approach contradicts the Biden Administration’s position that it is United States policy to not tolerate domestic monopolization as the answer to the rising power of foreign monopolies.³⁶

In conclusion, we thank you for your willingness to grapple with the problem of concentrated corporate control of the resources required to produce AI and similar technologies. In light of the depth and complexity of the questions raised by the proposal questions, we encourage the Task Force to fundamentally reconsider whether the NAIRR is consistent with both the broader aims of the project to “democratize” AI research and the Biden Administration’s explicit commitment to challenging the concentrated power of the tech industry. We look forward to supporting the Task Force in developing alternative proposals.

Sincerely,

Amba Kak, Director of Global Policy, AI Now Institute

Brittany Smith, Policy Director, Data & Society

Dr. Sarah Myers West, Post-doctoral Researcher, AI Now Institute

Meredith Whittaker, Faculty Director, AI Now Institute

³⁴https://www.documentcloud.org/documents/21062393-national-security-letter-on-antitrust?responsive=1&title=1&utm_source=newsletter&utm_medium=email&utm_campaign=newsletter_axioslogin&stream=top

³⁵<https://www.theverge.com/2019/10/17/20919464/mark-zuckerberg-facebook-china-free-speech-georgetown-tiktok-byte-dance>; <https://techcrunch.com/2019/07/17/facebook-or-china/>

³⁶<https://www.whitehouse.gov/briefing-room/presidential-actions/2021/07/09/executive-order-on-promoting-competition-in-the-american-economy/>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

AI Redefined, Inc.

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



Name of the person(s) or organization(s) filing the comment:

AI Redefined, inc.

400, McGill St
Suite 300
Montreal, Quebec
Canada
H2Y 2G1

About AIR

AI Redefined (AIR) is a Montreal-based company that has built a new AI training approach where humans and machines learn continuously from each other to address sophisticated challenges in real time. AIR recently released Cogment, the world's first open-source framework to provide the means to design, train, and deploy complex intelligent ecosystems that mix humans and AI agents of various kinds. In Cogment's real-world and simulated environments, actors, which can be humans, AI agents, traditional algorithms, or even aggregates of actors, can build trust and explore context together. Cogment can support human-AI collaboration on tasks like pilot training, the operation of hybrid vehicle fleets, and crisis management decision making. AIR currently works with leading players in the European aerospace and defense industries.

1. *What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]*

A. *Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;*

N/A

B. *A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:*

i. *An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and*

ii. *A governance structure for the Research Resource, including oversight and decision-making authorities;*

N/A

C. *A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;*

N/A

D. *Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;*

Among the many capabilities required to address this particular point, we believe that an open-source orchestration platform enabling the access and combined use, in shared environments, of different AI agents, as well as human users, is of paramount importance for several reasons:

- The ability to easily compare and audit implementations of AI agents with:

- other AI implementations,
- human users
- other non-learning algorithms or heuristics;
- The ability to keep human oversight;
- The ability to provide context to AI agents through human expertise.

Our multi-agent & human-in-the-loop open source AI framework, Cogment, allows just that, and thanks to its microservice architecture, is both scalable and distributable out-of-the-box. As a result, researchers worldwide can build Cogment projects to design, test, train, deploy and use AI agents and allow access to their work to their teams, other researchers, or users, in any configuration they choose. Training environments designed for specific research or projects can be shared as well, and we are currently working on additional tools to facilitate the sharing of agents and environments even further.

We stand ready to help the AI community build intelligence ecosystems mixing humans, heterogeneous AI agents and other systems.

- E. *An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;*

Open-source software (OSS) quietly affects nearly every issue in AI policy, but it is largely absent from discussions around AI policy—policymakers need to more actively consider OSS’s role in AI.

OSS is software that is free to access, use, and change without restrictions, and plays a central role in the development and use of artificial intelligence (AI). Across open-source programming languages such as Python, R, C++, Java, Scala, Javascript, Julia, and others, there are thousands of implementations of machine learning algorithms. OSS frameworks for machine learning, including tidymodels in R and Scikit-learn in Python, have helped consolidate many diverse algorithms into a consistent

machine learning process and enabled far easier use for the everyday data scientist.

Cogment core framework is OSS, and so can be projects built on Cogment. We believe that remaining open source, future proof, tech agnostic and accessible will not only make Cogment a better framework, it will also foster better research, and mitigate phenomena such as the digital fracture, biased AI / datasets, or other contributing factors impeding the general public's understanding, acceptance, and use of AI.

- F. *An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;*

N/A

- G. *An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;*

N/A

- H. *A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; and*

N/A

- I. *Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.*

N/A

2. *Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?*

N/A

3. *How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those*

concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

N/A

- 4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?*

N/A

- 5. What role should public-private partnerships play in the NAIRR? What examples could be used as a model?*

N/A

- 6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?*

N/A

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

The Alexandria Archive Institute (Open Context)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Response to the Request for Information on a National Artificial Intelligence Research Resource - RFI 86 FR 39081

From: The Alexandria Archive Institute (Open Context) prepared by Paulina F. Przystupa, with Eric C. Kansa and Sarah Witcher Kansa

Submitted on: September 23, 2021

1-a. Options for the goals for the establishment and sustainment of a NAIRR and metrics for success should include a heavy emphasis in education, demonstrable benefit to communities whose data are used to build the curated data sets used by the resource, an incorporation of inclusive governance for the resource, and improvement to infrastructure that allows easier access to communities currently underserved by AI research.

These should be included because without an infrastructure of education that leads into the use of such a resource it will probably be underused and not serve all communities in an equitable way. This includes education in the resource itself but also in the creation of the resource that allows multiple publics to understand its overall importance so that they can better influence its use and use it for themselves.

A demonstrable benefit to the communities whose data are used is important because scientific research has a history of exploitation that a resource like this has the potential to reinforce. These communities should also help define “metrics for success” as it is their information that is used to create these models.

An incorporation of inclusive governance for the resource is important because understanding what data and models may be appropriate will vary cross-culturally. This is especially important with respect to communities that have experienced colonialism and other forms of inequality. Such communities must be included in the definition and evaluation of any "metrics for success".

An improvement to infrastructure that allows easier access to communities currently underserved by AI research is important because this will allow for a longer sustainment for the NAIRR. Infrastructure improvements, particularly in underserved AI research communities, is key to the sustainment of NAIRR because it will ensure that a wide variety of communities will continue to use the resource. Additionally, this access will open up as yet unexplored possible avenues for the utilization of such a resource. This includes things like cheaper and faster internet access and computing technology.

1-b-ii. Options for roadmap element 1.B.ii should consider regarding a plan for ownership and administration for NAIR and specifically the governance structure for the NAIR, including oversight and decision-making authorities, are ones that should forefront issues of Indigenous Data Sovereignty and, more generally, a framework of

Inclusive Governance. This is important because it will acknowledge that not all data curated by the resource nor produced by it has the same potential for impact or influence in all communities. It should acknowledge that diverse data should come with diverse structures and cultural requirements for its governance. However, where data pertains to communities not historically minoritized, datasets and curated should strive to follow ethical open data principles and be in the public domain.

Furthermore, the roadmap should also consider a plan for ownership and administration of NAIRR including a governance structure, with oversight and decision-making authorities, that informs subjects who were part of research that created data training sets about when information about them is available or being used. This is important because it will ensure that the research developed out of NAIRR maintains a reciprocal and responsible relationship with the communities who are directly or indirectly incorporated in the work.

Additionally, the roadmap should also consider a plan for ownership and administration of NAIRR including a governance structure, with oversight and decision-making authorities, that includes the possibility for the retirement of data sets. A right to forget should be incorporated into this work and is important because not all information is necessarily collected to be used in perpetuity. In addition, our understanding of privacy may change where information about individuals, places, or other things included in this resource may no longer be considered to have been ethically obtained or relevant to the resource.

Lastly, the roadmap should also consider a plan for ownership and administration of NAIRR including a governance structure, with oversight and decision-making authorities, that acknowledges historical inequalities in the creation and construction of resources and training data sets with regards to this research. Specifically, not all groups have historically been treated with the same fairness when data sets are generated about them. Therefore, the governance structure should actively ensure that equitable application and data use exist. This is important because it will build trust in the resource amongst a wide variety of communities that could benefit from this research by acknowledging histories of differential access to such materials.

1-c Options for the goals for a model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources should include a system similar to those of institutional review boards for human-subjects research in the sciences.

The goals for this model should include a focus on ethical and equitable governance and oversight that allow those in charge to establish ethical strategic directions, programmatic decisions, and allocate resources. This is important because it foregrounds the ways that such a resource can explore inequalities in use, education, and maintenance as well explore who the resource is really serving.

Additionally, the goals for the models for governance and oversight should utilize perspectives from outside direct AI research. This is important because those in the digital humanities and ethics explore the direct and indirect impacts of the resource that will aid in the creation of the decisions.

1-e Options for an assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource should include funding education in ethics and all fields that the resource intends to use data sets from (including education in racial and gender equity, fairness, bias, civil rights, transparency, and accountability), inclusive governance, and lastly, education in the technologies that underpin AI research that begin at the elementary level.

A lack of current funding in the above listed fields is a barrier to the dissemination of NAIRR products because it means a limited community can utilize the resource. A recommended solution to this is to fund education in these fields. Early education in these fields, additionally, will cultivate new and innovative ideas for the use of the resource that will better address problems currently facing American communities.

The lack of inclusive governance of adjacent resources is a barrier to the dissemination of NAIRR products because it means that the products will probably reinforce existing hierarchies in AI knowledge production. A recommended solution to this barrier is to incorporate practices of Inclusive Governance that draw from multiple cultural ways of understanding how knowledge and research should be conducted and the rights of certain communities to withhold information.

Lastly, the lack of education in the technologies that underpin AI research is a barrier to the dissemination of NAIRR products because it means that many people who may be interested in AI do not have the prerequisite skills to undertake this research if they approach it later in their educational careers. A recommended solution to this barrier is to add education in AI as part of elementary education. This should include education in the previously listed fields but also in the technologies that underpin AI so that a holistic approach to the use and importance of the resource is encouraged at the beginning of one's educational career.

1-g Options for an assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research should include an examination and restriction of the data private companies are able to retain or collect on individuals whose data sets might be incorporated and the right to forget or not preserve all material that goes into this resource.

This is important because our private companies are already able to retain a lot of information about users without accountability and to have a government agency who

will directly benefit from information by or about particular communities do that as well would not support civil rights and civil liberties.

The right to forget or have curated data sets be removed is also important because it acknowledges that not all collected data are necessarily useful forever. Particularly, information that pertains to individuals should have a limited shelf life to preserve their long-term rights to privacy.

2. The capabilities and services provided through the NAIRR that should be prioritized are educational tools and services, including provision of curated data sets, and secure access control.

3. The NAIRR and its components can reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability, by being built with the aid of ethical and responsible creators. This particularly means the incorporation from the beginning of ethicists as well as individuals outside of current AI research who do the work on racial and gender equity, fairness, bias, civil rights, transparency, and accountability. Professionals in those fields will be the ones who can speak expertly on the issues most pertinent to AI research that will make the tool actually responsible and ethical.

It can also do this by having the program under an agency that is held more accountable by the public than the current funding and oversight agency. For example, many communities, especially communities with histories of disenfranchisement, would not trust Federal agencies with law enforcement or national security responsibilities to prioritize ethical and responsible research. Federal agencies with missions that directly promote human welfare and scientific progress would have more trust.

Lastly, NAIRR can reinforce those principles by focusing on the serious vetting of any and all training datasets. All training datasets must go through rigorous ethical and responsible scrutiny to ensure that they are not constructing bias that could have severe negative impacts on the previously stated topics. This vetting requires the incorporation of individuals from a variety of backgrounds so that the lifecycle of data sets is explored in terms of creation, management, dissemination, and incorporation to evaluate all possible sources of harm that could come from that data set. To do this requires many voices to be incorporated into the process.

4. There are a number of building blocks that already exist for NAIRR in terms of government, academic, or private-sector activities, resources, and services. The first suggestion would be for the OTSP and NSF to examine their currently funded projects to see which ones include artificial intelligence research, open data infrastructures, and

studies on related topics in racial and gender equity, fairness, bias, civil rights, transparency, and accountability. These already government-funded projects should provide a significant resource for the creation of curated data sets as well as investigations into secure access control, and the other items listed in 1.D.

Our organization's suggested building blocks primarily focus on organizations involved in the creation of educational tools and services as well as in the provision of curated data sets. Open Context, a project of the Alexandria Archive Institute, provides both educational tools and services as well as curated data sets that could be incorporated or built from for NAIRR. Other possible sources of open data sets include the iSamples project, tDAR, and WholeTale.

5. The role that public-private partnerships play in the NAIRR should be limited to private partnerships that focus on racial and gender equity, fairness, bias, civil rights, transparency, and accountability. Specifically, public-private partnerships should be sought with private entities that focus on the ethics and social justice outcomes that NAIRR may be interested in cultivating. Public-private partnerships that focus on the ethical and just outcomes for the research will be better able to utilize the product as well as ensure the construction of a resource that acts in an ethical fashion.

6. We see the see limitations in the ability of the NAIRR to democratize access to AI R&D in its affiliation, guidance, and funding from the DOD; if it does not incorporate questions of racial and gender equity, fairness, bias, civil rights, transparency, and accountability from the beginning; if scaffolding programs that include improvements to elementary education in AI related fields including racial and gender equity, fairness, bias, civil rights, transparency, and accountability studies are undervalued; and if it doesn't actively seek out methods to ethically democratize access.

To better promote democratization of AI R&D, we recommend that NAIRR should be entirely separate from National security and law enforcement agencies. Administrative control or funding from these agencies would undermine the credibility of NAIRR attempts to develop and promote responsible use of AI methods across the broader civil society. Instead, national security and law enforcement agencies should seek best practice, governance, and ethical guidance from initiatives supported by NAIRR, but NAIRR itself should operate independently from national security or law enforcement.

The limitation to democratizing access to AI R&D due to a lack of the incorporation of questions of racial and gender equity, fairness, bias, civil rights, transparency, and accountability from the beginning occurs because without an investigation of these issues equitable democratization is impossible because it will not be fully defined. Without identifying groups and areas of interest that lack access or have unequal access, it will not be possible to make that access equitable amongst all

groups interested or actively recruited into use of the resource. this can be overcome by incorporating these questions from the beginning of the tool's development process

The limitation to democratizing access to AI R&D due to a lack of scaffolding programs that begin with elementary education means that only certain kinds of researchers will have interest and access. The education required to develop AI technologies and ask ethical questions that AI could help solve, should start at the elementary level as early education often becomes a predictor for later outcomes. The ability to get funding and access to AI research is already an undemocratic system. Therefore, for NAIRR to actually change that, it would need to work to undercut the existing system of tiered access to AI education beginning at the elementary level. This could be solved by expanding elementary education funding in underperforming schools that support STEM, STEAM, and humanities education more generally, that will work together to provide students with the knowledge to succeed and utilize NAIRR.

The limitation to democratizing access to AI R&D due to a lack of inclusive governance could lead to democratization without justice. This could result in the continued marginalization of groups in the Nation. For example, Indigenous peoples make up a small percentage of the population; however, it is possible that data sets built from research that exploited Indigenous knowledge or peoples could be integrated with this resource and continue to do harm by not acknowledging Indigenous Data Sovereignty. In that case, it would be a democratized resource, but its impact could be disproportionate on certain groups. However, if these datasets do not include data about minoritized groups in the United States, because democratization can come into conflict with sovereignty and the rights of minoritized people, this resource will be inequitable, serving or creating research about only some individuals in the Nation. Democratized access that acknowledges the need for tiered access to AI R&D will provide a most just method for the resource.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Amazon Web Services

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

1 Introduction

In an effort to democratize the infrastructure to fuel artificial intelligence (AI) research and development (R&D), Congress established the National Artificial Intelligence Research Resource (NAIRR) Task Force (Task Force). The Task Force has been mandated by Sec 5106 of the National Artificial Intelligence Initiative Act of 2020 to “develop a coordinated roadmap and implementation plan for creating and sustaining a National Artificial Intelligence Research Resource.” The NAIRR would establish a shared infrastructure that, per the RFI, “would provide Artificial Intelligence (AI) researchers and students across scientific disciplines with access to computational resources, high-quality data, educational tools, and user support.” To support this initiative, Amazon Web Services, Inc. (AWS) is responding to the RFI to offer our perspective as a leading cloud and AI provider with years of experience building AI solutions for government, nonprofit, and educational organizations around the world.

AWS and CloudBank

AWS supports groundbreaking research by working with higher education institutions and national and international research agencies to advance research and promote collaboration using the AWS Cloud. Over the past decade, we have supported key initiatives, such as the NSF-funded [CloudBank](#), for data and compute-intensive research and education efforts using the cloud.

AI and related technologies, including machine learning (ML) and deep learning, enable government agencies to transform how they operate and improve citizen services. We know this from our experience helping the federal government and research institutions create impactful AI solutions faster. As a leading cloud service provider, AWS’s compute, storage, AI/ML, and data analytics services can form the backbone of NAIRR’s shared research infrastructure, helping the Task Force meet the requirements in the [National Artificial Intelligence Initiative Act of 2020](#) and as directed by Congress.

AWS seeks to continue supporting research and education efforts and bridging the digital divide through initiatives like the NAIRR. Our AI services offer capabilities to support the Task Force in developing a roadmap for “expanding access to critical resources and educational tools that will spur AI innovation and economic prosperity nationwide,” as stated in The White House Office of Science and Technology Planning (OSTP) press release, *The Biden Administration Launches the National Artificial Intelligence Research Resource Task Force*. Additionally, our suite of AI services enables the Task Force to carry out the activities described in the OSTP press release and in the National Artificial Intelligence Initiative Act.

Section 2 includes our responses to the RFI questions and our recommendations for the Task Force to consider during development of the NAIRR roadmap and implementation plan.

2 Responses to the RFI Questions

Our responses to the RFI questions provide our recommendations for establishing and sustaining the NAIRR and meeting its objectives as a shared research infrastructure.

2.1 Question 1: Implementation Roadmap Elements

As outlined in Section 5106(b) of Public Law 116-283, the implementation roadmap developed by the Task Force should include the following elements discussed in the subsections below.

2.1.1 Roadmap Element A: Goals and Success Metrics for Establishing and Sustaining the NAIRR

To further the National AI Initiative and create a NAIRR that is sustainable for the future, the NAIRR should establish goals and associated metrics that clearly demonstrate the efficacy of the program. Good metrics will enable data-driven decision making related to future funding and the efficacy of public and private sector investment in the NAIRR. While there are many goals the Task Force can consider, we recommend emphasizing goals that increase the benefits of AI technology to underrepresented and underserved communities and that strengthen our national AI workforce (such as the suggested goals and metrics below). As AI technology is introduced more widely to the U.S. population through the NAIRR and other initiatives, the demand for AI education and a developed workforce will continue to grow.

Suggested Goal 1: Increase the extent to which AI technology benefits underrepresented and underserved communities. AI technology has the potential to improve myriad critical products and services for all segments of society (e.g., legal, healthcare, financial services, etc.). The NAIRR can track the use cases and communities served by the technologies developed with its resources and adjust the allocation of its resources over time to ensure that underrepresented and underserved communities benefit from the technologies developed. The NAIRR will need input from diverse stakeholders to identify the communities and use cases that it should prioritize and to decide what qualitative or quantitative measures of “benefit” make the most sense. For example, for a small number of specific demographic and healthcare use cases each year, the NAIRR could (a) inventory the publicly available AI applications that address the use case, (b) assess the size of the gap between what a National Institutes of Health (NIH) research panel suggests is possible and the inventory, (c) prioritize allocating computing resources to classwork and research that closes the gap, and (d) measure the closure of the gap.

Suggested Goal 2: Increase the size and skills of the national AI workforce. To take full advantage of the opportunities that AI presents, we need to increase the number of people qualified to fill the new technical and non-technical jobs being created by the AI boom. While there are already programs to increase the number of STEM graduates, within AI, STEM roles constitute only a portion of the new jobs that will become available. For example, there will be new AI-related jobs in product design, dataset construction, policy analysis, and business leadership. All of these jobs will require a solid understanding of AI principles and some familiarity with how real systems work. The NAIRR can (a) review the availability of AI curricula and research programs across secondary and post-secondary schools in the United States, (b) estimate the production rate of skilled technical and non-technical workers from existing programs, (c) sponsor the development of new curricula and research programs that use NAIRR compute, storage, and other resources to accelerate skill development, (d) help schools without existing programs or with underdeveloped programs expand their enrollment of students via the adoption of the new NAIRR curricula, and (e) track the resulting improvement in the national AI workforce.

2.1.2 Roadmap Element B: NAIRR Ownership and Administration

NSF, or NSF-led National AI Research Institutes, may be best suited to operate the NAIRR's infrastructure. NSF has existing data computing and shared infrastructure through partnerships with cloud service providers. This infrastructure can support the NAIRR's implementation, deployment, and administration. NSF can leverage this existing infrastructure to provide the NAIRR as a resource to universities and colleges, nonprofits, and other educational or research institutions focused on AI. This collaboration can help allocate resources and knowledge to support the NAIRR's ownership and administration of its objectives—whether through infrastructure, training services, or programs—and evaluation of responsible AI research. Recommendations for the NAIRR's governance structure are described in **Section 2.1.3**.

2.1.3 Roadmap Element C: Governance and Oversight Strategy Model

NSF, with its own oversight and administration, is well equipped to support programmatic decisions and allocate resources to researchers and students. Supplementary to NSF, we recommend creating an advisory board composed of advocates for underrepresented and underserved communities, professional organizations dedicated to economic development, and AI experts from the private sector and federal agencies. The purpose of this board is to gather individuals who have experience with critical technology initiatives and making AI more accessible, enhancing online adult learning programs to upskill the workforce, and supporting underrepresented students in elementary through post-secondary education. This advisory board can help set goals for the NAIRR, track progress against those goals, and allocate resources to achieve those goals.

We recommend an iterative approach to the governance and oversight strategy. Programmatic decisions should initially be focused on tangible outcomes and agile iteration of minimum viable products (MVPs), such as a rapid prototype of an AI model or a pilot education program. Data produced from these MVPs can inform how to improve a prototype or pilot before scaling it to be an official component of the NAIRR. By starting with small-scale, agile iteration of the NAIRR's critical components rather than larger-scale investments, the NAIRR could achieve early success and accelerate buy-in from the government, research institutions, and the private sector.

2.1.4 Roadmap Element D: Creating and Maintaining Shared Computing Infrastructure

Research institutions like NSF have decades of experience creating and maintaining shared computing infrastructure. Together with modern infrastructure providers, including large-scale cloud service providers, the NAIRR can leverage the expertise of research institutions with previous experience establishing and maintaining the infrastructure. The core capabilities and resources of the NAIRR should include processes to communicate the NAIRR's mission to students and researchers, enroll students and researchers in the NAIRR programs, grant access to the NAIRR resources,

Using AWS to simplify access to cloud resources for researchers

Using AWS, researchers have created the largest high-performance cluster in the cloud to study natural language processing by using over 1.1 million virtual CPUs running in a single AWS Region, demonstrating the scalability and elasticity of the infrastructure.

track usage and results of using the NAIRR, and contract with vendors to provide underlying software systems.

AWS offers several capabilities to support creating and maintaining a secure shared computing infrastructure like the NAIRR and to give researchers access to advanced computing resources,

training, and curated datasets. For example, the [AWS Global Cloud Infrastructure](#) enables the NAIRR to deploy application workloads across the globe in a single click or build and deploy specific applications closer to end users with single-digit millisecond latency so that researchers can quickly analyze massive data pipelines, store petabytes of data, and advance research using transformative technologies like AI/ML. Additionally, in support of initiatives like the NAIRR, AWS provides curated public datasets available to researchers through programs like the [Registry of Open Data on AWS](#).

2.1.5 Roadmap Element E: Assessment of and Solutions to Use of Government Datasets

Access to data is crucial for research in AI, including access to high-quality, publicly available data and proprietary datasets. The NAIRR therefore must make it easy to find, subscribe to, and use valuable government datasets in the cloud. Currently, use of government datasets can be a complex issue since it is often difficult to make government datasets easily available and accessible to researchers because of concerns about data protection and privacy or because agencies have voluminous data stored within siloed and disparate data sources.

To help address these challenges, the NAIRR should support continued research and enable piloting of innovative data sharing mechanisms that could help reduce obstacles to the use of government datasets. For example, the [Global Partnership on AI](#) is pursuing comprehensive research on data trusts, or a legal framework for managing shared data, as part of its mandate; the Task Force should consider a partnership to further support this work.

Additionally, solutions to the challenges of using government datasets should include the ability to quickly find datasets in a shared repository, securely access datasets from the cloud, and easily characterize datasets. For example, the NAIRR can use research and technical computing resources like the [AWS Data Exchange](#), [Registry of Open Data on AWS](#), [Amazon Sustainability Data Initiative \(ASDI\)](#), and [AWS Open Data Sponsorship Program](#) to support the computational demands of training and collecting, storing, and sharing data at scale.

2.1.6 Roadmap Element F: Assessment of Security Requirements

The key security requirements for the NAIRR are (1) securing the datasets, models, and other work products of NAIRR users, (2) controlling access to the computing resources within the shared infrastructure (e.g., ensuring each user consumes only the resources allocated to them), and (3) limiting access to data within the NAIRR's administrative systems, which may overlay or interact with the shared infrastructure (e.g., ensuring shared resources are, in fact, shared and not exploited by subgroups of users).

Additionally, the NAIRR should consider trade-offs between ease of public access to data and users' willingness to share data. For example, to avoid restrictive policies for sensitive datasets,

the NAIRR could consider only hosting datasets that are publicly available, which may encourage dataset owners to make their data more broadly available. These public datasets would help reduce barriers to accessing proprietary data that would otherwise not be shared with the NAIRR's users. Requiring publicly available data will increase the number of people who can access the data, but the NAIRR should be mindful that this requirement may decrease the amount of data available to users.

2.1.7 Roadmap Element G: Assessment of Privacy and Civil Rights and Liberties Requirements

Developing responsible AI solutions is a process involving inputs and discussions with key stakeholders during all stages of the ML lifecycle, including AI developers and engineers, AI-focused policy experts, and end users and communities. The use of AI and ML technology, such as facial recognition, must comply with all laws, including laws that protect civil rights. There should be no ambiguity that existing laws (e.g., the Civil Rights Act of 1964 and the Fourth Amendment to the U.S. Constitution) apply to and may restrict the use of AI/ML technology in some circumstances.

Standards that establish clear benchmarks and testing methodologies are a proven way to address design issues in software, including fairness in AI. AWS supports current initiatives to develop independent standards for AI/ML with the International Organization for Standardization (ISO) and National Institute of Standards and Technology (NIST). The Task Force should adopt and support the development of such standards and guidelines while driving responsible research behavior and improvement of research outcomes.

AWS has engaged with NSF, NIST, and other stakeholders to offer direct assistance to this effort. NSF and AWS, for example, are partnering to support computational research on fairness in AI (see [NSF Program on Fairness in Artificial Intelligence in Collaboration with Amazon](#)). We also support efforts by members of the academic community to establish independent and trusted criteria, benchmarks, and evaluation protocols around AI/ML services (e.g., AWS Machine Learning Research grants invite researchers to apply for funding in various research areas, including ML topics such as fairness, privacy, and explainability in AI).

2.1.8 Roadmap Element H: Plan for Sustaining the NAIRR through Federal Funding and Private Partnerships

The U.S. government plays a significant role in funding long-term, critical research initiatives. The current administration is working on an AI policy will provide for continued investment in AI research to sustain U.S. technological leadership. Federal investments in AI R&D are imperative. For example, the [American Jobs Plan](#) proposes a \$180 billion investment to advance U.S. leadership in critical technologies, upgrade research infrastructure, and remove inequities in research and

Partnering with educational institutions to increase access to computer science education for children and young adults

[Amazon Future Engineer](#) is a comprehensive childhood-to-career program aimed at increasing access to computer science education for children and young adults from underserved and underrepresented communities.

careers in emerging technologies. Additionally, Division E of the 2020 National Defense

Authorization Act allotted significant funding to position the U.S. as a global leader in AI. As the NAIRR supports many of the objectives of these plans, funding for the NAIRR should be included in any current significant AI funding initiatives. Ultimately, the NAIRR should request long-term funding to reflect the government's investment in sustaining long-term AI/ML goals domestically and internationally.

Collaborations with the private sector, post-secondary institutions, and non-governmental organizations can also help sustain the NAIRR through access to talent, new ideas, and innovative approaches. AWS, for example, offers technical resources, mentorship, and networking opportunities to help build a future of technology that is inclusive, diverse, and accessible. With programs such as [We Power Tech](#), [AWS Educate](#), and [Amazon Future Engineer](#), AWS increases the number of underrepresented technologists participating in the innovation economy through educational content, partnerships, and programs. The NAIRR can use these AWS examples and similar programs to connect with a national network of future AI students and leaders.

Resources for researchers and institutions to work remotely

The rapidly changing and dynamic global health situation has impacted the lives of many people, including researchers at universities and institutions worldwide. Many academic institutions are migrating to remote operations. Researchers are processing data, collaborating online, and maintaining labs remotely. Amazon and AWS are responding to these events in support of our communities and deploying resources and technology to enable remote learning and working from home. Visit our [AWS Public Sector Blog](#) for more information.

2.1.9 Roadmap Element I: Parameters for Establishing and Sustaining the NAIRR and Agency Roles and Responsibilities

Parameters for establishing and sustaining the NAIRR should primarily follow standards and best practices already in place for public sector AI initiatives. For example, NSF and NSF-led AI Research Institutes and NIST have technical standards that promote innovation and public trust in systems that use AI, including standards for AI data, performance, and governance. For additional industry recommendations about agency roles and responsibilities, refer to the considerations in the preceding roadmap elements.

2.2 Question 2: Capabilities and Services Prioritization

The Task Force should prioritize capabilities and services that focus on investing in and managing IT, AI/ML programs, and research projects aligned with the NAIRR's goals and objectives. Priorities should include facilitating access to advanced cloud computing resources and curated datasets and delivering educational tools and programs to advance U.S. leadership in critical technologies.

Simplifying Access to Cloud Computing Resources for Research and Education Efforts

Having the right computing resources is necessary to sustain the NAIRR. The AWS Cloud and other public clouds enable researchers to quickly and affordably access the latest versions of

many resources that may otherwise be difficult to obtain. This level of availability promotes deployment of new ideas or services that otherwise could have taken months or years to achieve.

Further, AWS pre-trained AI services can provide ready-made intelligence for the NAIRR's applications and workflows.

AI services easily integrate with applications to address common use cases such as extracting text from unstructured documents, automating document and image analysis, adding natural language search capabilities, and building accurate forecasting models.

Bridging the Digital Divide

Training and education are fundamental to the NAIRR's effort to democratize access to AI/ML technology. This should include teaching researchers and students how to approach and process complex workloads by providing scalable and secure compute, storage, and database capabilities to accelerate time to science.

There are many commercially available training and education programs available to the NAIRR and NAIRR users. [AWS Training and Certification](#) has curated a list of no-cost, on-demand online courses tailored to researchers' needs. These online courses are available at any time to help users learn new cloud skills. For example, [Research Learning Pathway: Foundational Services](#) is a no-cost online AWS training pathway for researchers and research professionals who want to become more proficient in optimizing research on AWS. AWS also has low- or no-cost training and education resources for learners. Cloud-based educational programs like [AWS Academy](#), [AWS Educate](#), and [AWS re/Start](#) are designed to reach diverse communities and individuals, including underserved and underrepresented populations.

Reaching diverse communities and individuals with varying skill levels

AWS Academy and AWS Educate partnered with the Los Angeles Community College District to offer Cloud Computing Certificate training at 19 Los Angeles County community colleges and their sister high schools.

2.3 Question 3: Ethical and Responsible Research Development

With the increased use of AI in everyday life, creating fairness in AI systems is critical. Amazon collaborates with NSF and other research institutions to award research grants in areas such as ensuring fairness in AI algorithms (and the systems that incorporate them), using AI to promote equity in society, and developing principles for human interaction with AI-based systems. The agency administering the NAIRR and its components can promote principles of responsible AI R&D by implementing the following recommendations.

Recommendation 1: Promote diversity and inclusion within NAIRR staff, within NAIRR student, educator, and researcher populations, and within the AI development process. The NAIRR can improve the outcomes of using AI technology by including a variety of perspectives (representing different backgrounds, skillsets, beliefs, and life experiences) in the design, development, and operation of its policies and practices; in the student, educator, and research populations that it supports; and in the design and testing of AI applications built using NAIRR resources. For example, the NAIRR will need to set policies that determine how much compute

and storage to allocate to each student, class, educational institution, and research team. These policies will have downstream impacts on the economy and so should be influenced by a broad spectrum of stakeholders (e.g., end users, technologists, academics, industry experts, lawyers, parents). As another example, the NAIRR can incorporate mechanisms to help researchers engage with diverse groups of citizens and other end users during the design and assessment of AI technologies.

Recommendation 2: Fund the datasets, research, and education needed to improve AI fairness, explainability, privacy, and transparency. First, since the lack of high-quality training and test datasets often impedes AI research, the NAIRR should fund the creation and acquisition of datasets that represent all demographics (for fairness research), that capture complex causal relationships between data attributes (for explainability research), and that otherwise enable privacy and transparency research. Second, the NAIRR should sponsor research, possibly funded in partnership with other public and private sector organizations, focused on the following areas: technical solutions for fairness, explainability, privacy, and transparency; best practices for educating a national AI workforce; and best practices in building and using AI applications. These research programs should be executed using NAIRR compute, storage, and datasets. Finally, the NAIRR should fund educational programs to help the public better understand the nature and capabilities of AI systems and increase the influence of the public on the design and use of AI applications. These educational programs should cover important concepts like the predictive nature of ML, confidence indicators and thresholds, the importance of human involvement and human review, capabilities and limitations of AI, recommended or prohibited uses, and best practices.

2.4 Question 4: Building Blocks for the NAIRR

By working with research institutes and cloud service providers, the NAIRR gains access to existing programs and services that can be used as building blocks to customize engagement across the public sector. The agency administering the NAIRR can use these building blocks “out of the box” to move with increased speed. For example, in addition to providing cloud infrastructure services (**Section 2.2**), we also invest in successful public-private partnerships across government and academia, which are detailed in **Section 2.5**. The NAIRR can incorporate these partnerships to upskill a community of AI users in the future.

2.5 Question 5: Public-Private Partnerships

Public-private partnerships between education, industry, and policymakers should play a critical role in the NAIRR because they can speed up access to educational resources and stackable credentials mapped to in-demand tech jobs. AWS has implemented public-private partnerships that have successfully trained workforces, brought emerging technology to under-resourced communities, and provided educators with cloud-focused curricula. Our programs can inform a future model for the NAIRR to align public programs of study with access to cloud technology.

We have teamed with educational institutions across the U.S. to offer cloud skills education as part of matriculating and non-matriculating degree programs at scale. We also support the workforce and economic development efforts of state and local governments through cloud-focused

programs that upskill people with varying educational backgrounds. Our programs, a few of which we detail below, offer instructional resources, curriculum alignment to industry demand, faculty development, career support services, certifications, academic services, and connections to employers looking to hire cloud-skilled talent.

Connecting Academia and AWS

AWS has public-private partnerships with NSF, researchers, and institutions to advance emerging technology initiatives like the NAIRR. Through [AWS Academy](#), we provide higher education institutions with a free, ready-to-teach cloud computing curriculum that prepares university students to pursue industry-recognized certifications and in-demand cloud jobs. This curriculum includes AWS Academy Machine Learning Foundations, which can help certify educators to teach AI/ML concepts.

NOVA and AWS Announce Cloud Computing Degree

Through a public-private partnership, AWS Educate, AWS Academy, and Northern Virginia Community College (NOVA) launched a [cloud computing specialization](#) as part of NOVA's Information Systems Technology Associate of Applied Science degree in Fall 2018, making NOVA one of the first community colleges to offer a cloud computing degree.

Free Hands-on Training for Any Individual

As part of our efforts to make the cloud accessible to all, [AWS Educate](#) provides an online suite of free hands-on labs and personal study resources for any individual interested in learning about cloud fundamentals on AWS, including a Machine Learning Pathway. This pathway can help individuals discover how to use the AWS AI suite to solve real-world problems. By the end of the course, individuals gain the prerequisite skills and competencies necessary to be successful as an ML scientist.

Accelerating Research through Collaboration

AWS collaborated with NSF and the NIH to create the [AWS Research Initiatives \(ARI\)](#) program, which supports groundbreaking research in the fields of computer science, biomedical engineering, and information science. ARI provides support to public sector organizations through access to AWS scientists, research workshops led jointly by AWS and NSF/NIH, AWS research credits that can be applied to AWS Cloud service usage, and awards consisting of federal funds and AWS resources.

2.6 Question 6: Limitations to Democratizing Access to AI/ML

Limitations to the democratization of AI R&D are multifaceted, often driven by technical, socioeconomic, or educational factors. For example, underserved communities or smaller organizations may experience challenges accessing AI/ML because of large upfront costs. However, cloud-based solutions are helping to break down barriers to entry by bringing computational capabilities within reach of more public sector citizens, regardless of prior technology expertise, funding, or geographic location. Based on our own efforts to democratize AI/ML, we provide detail below on the limitations we foresee for the NAIRR in addition to our recommendations for overcoming these limitations.

Privacy and Security Limitations

As with any technological advancement, expanded access may lead to expanded opportunities for risk, such as bad actors that would be a threat to cybersecurity. As the NAIRR expands access to AI resources, the privacy and security controls associated with AI services must scale alongside usage. The NAIRR should leverage a robust underlying infrastructure with built-in privacy and security components—such as commercial cloud—to proactively protect users from cybersecurity attacks.

Educational Limitations

Lack of instructor expertise is a primary challenge that public and higher education institutions face when enabling AI democratization through education. These institutions need time, resources, and tools to upskill their faculty and keep their programs current with rapidly evolving AI technology. To address knowledge gaps as structured degrees and a standardized, national curriculum are more widely adopted for this technology, AWS provides U.S. educators with free professional development sessions, access to free self-paced learning resources, and free instructor certification exams. The NAIRR can promote resources like AWS Academy to help educators and students gain access to cloud skills.

Financial Limitations

Financial constraints may present obstacles to the democratization of AI/ML technology. To overcome them, we recommend the agency administering the NAIRR engage with industry to explore possible funding models that would enable private industries to contribute service credits or additional funds to initiatives related to AI/ML research and accessible education. Additionally, the agency administering the NAIRR should seek government seed funding to substantiate the government's investment in the outcomes of the NAIRR.

Capacity Limitations

AI R&D requires tremendous data storage and high-speed, high-volume connectivity to enable the accessibility of development, training, and use of AI technology. Further, fiber-optic broadband connections—and the skilled technical workforce needed to construct, maintain, upgrade, and secure them—are essential to providing access to the NAIRR. The cloud provides the compute power and virtually unlimited scalability required to sustain AI R&D and access to the NAIRR. For example, AWS uses 100 Gbps network bandwidth to provide high-speed access to those in rural communities so that they have more opportunities to expand their tech skills and career development.

3 Conclusion

As the need for advancement in AI grows, AWS stands as an industry leader with a robust, reliable, and scalable cloud infrastructure. We are dedicated to supporting the federal government in its AI initiatives to advance AI R&D. We are also committed to building a diverse and highly skilled workforce that can close the AI talent gap. To discuss our response to the RFI or the NAIRR initiative in general, please feel free to reach out for more information. We look forward to discussing this initiative with OSTP and NSF in more detail.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

American Civil Liberties Union (ACLU)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

RE: Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource (86 Fed. Reg. 39081)

I. Introduction

While there may be value in creating a shared resource where American researchers can build and develop artificial intelligence (“AI”) tools for many different applications, it would be a mistake to assume that investment in AI is inherently beneficial. The current framing of this resource takes the technochauvinist¹ view that the most pressing problem with AI is its inaccessibility and that the democratization of AI will inevitably lead to positive returns for the United States. But expansions of AI that lack an express, ongoing focus on how AI development will affect civil rights and civil liberties will invariably lead to technologies that threaten these important protections. Applications of AI are already used in decisions that fundamentally impact people’s lives, ranging from who gets held in jail to who gets a job, loan, or insurance. The discrimination in these uses of AI continues unabated. Regulation and legislation have not kept pace, failing to identify and root out AI bias or halt applications of AI where the risk of discrimination is too great. People are increasingly becoming aware of the ways AI and algorithmic decision-making systems may already be negatively affecting them, causing a degree of “techlash.”² Current questions around AI abound: when does prediction lead to discrimination; when do the goals of increased safety and efficiency creep into unwelcome surveillance; and in what ways does current AI innovation trap us in past assumptions, dooming the future to repeat the past? These questions must be grappled with and addressed from the outset as part of any initiative to encourage AI development.



National Office
125 Broad Street
18th Floor
New York, NY 10004
aclu.org

Deborah N. Archer
President

Anthony D. Romero
Executive Director

II. Lack of Civil Society Representation on the Task Force and Ways the Task Force Can Constructively Engage Going Forward

Harms to the individuals and communities upon whom AI is ultimately deployed can originate at every stage of the AI lifecycle—in problem and domain select

¹ Meredith Broussard, *Artificial Unintelligence: How Computers Misunderstand the World* 7–8 (2018) (“Technochauvinism is the belief that tech is always the solution.”).

² See *The Techlash Against Amazon, Facebook, and Google—and What They Can Do*, *The Economist* (Jan. 20, 2018), <https://www.economist.com/briefing/2018/01/20/the-techlash-against-amazon-facebook-and-google-and-what-they-can-do> [https://perma.cc/Z8YK-D5EL].

in data collection and selection, and in ideation, development, and deployment of AI. And these harms occur in critical areas impacting individual freedom as well as social and economic opportunity. Myriad forms of bias and harm have already been identified in the collection and use of data and deployment of AI in the criminal legal system,³ housing,⁴ employment,⁵ credit,⁶ and education.⁷

The NAIRR Task Force has a critical role to play in determining the path that AI development will take, and robust consideration of these potential harms must be front and center in the Task Force’s work. As discussed in more detail below, this includes decisions about ethical standards for whether or not to make certain data sets available in the first place, whether or not to support the development of particular AI tools, and how to identify and prevent uses of the NAIRR that would create unacceptable harms to individuals and communities. Further, there must be transparency and accountability in decision-making around what datasets are made available and the technologies that are ultimately deployed through use of those datasets.⁸ Additionally, in setting “parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities,”⁹ the NAIRR Task Force should (1) create specific and concrete agency goals, responsibilities, and milestones regarding substantive discussions to address a variety of bias, equity, transparency, accountability, and civil rights/liberties issues (“harms”) that may arise in the development, execution and maintenance of the NAIRR and (2) establish concrete ways to meaningfully identify and address these harms in the development, execution, and maintenance of the NAIRR.

The primary impediment we see to creation of a NAIRR that consistently and appropriately reinforces principles of ethical and responsible research, development, and deployment of AI is

³ Michelle Bao et al., *It’s COMPASlicated: The Messy Relationship between RAI Datasets and Algorithmic Fairness Benchmarks* (June 10, 2021), <https://arxiv.org/abs/2106.05498> [https://perma.cc/H44H-QBAG].

⁴ Letter from ACLU et al. on Addressing Technology’s Role in Housing Discrimination to U.S. Dep’t of Hous. & Urb. Dev. et al. (July 13, 2021), <https://www.aclu.org/letter/coalition-memo-re-addressing-technologys-role-housing-discrimination> [https://perma.cc/52RV-PNBF]; Patrick Sisson, *Housing Discrimination Goes High Tech, Curbed* (Dec. 17, 2019), <https://archive.curbed.com/2019/12/17/21026311/mortgage-apartment-housing-algorithm-discrimination> [https://perma.cc/M3XH-B3PD].

⁵ Letter from ACLU et al. on Addressing Technology’s Role in Hiring Discrimination to U.S. Equal Emp. Opportunity Comm’n et al. (July 13, 2021), <https://www.aclu.org/letter/coalition-memo-addressing-technologys-role-hiring-discrimination> [https://perma.cc/W39Q-SQKT]; Ctr. for Democracy & Tech., *Algorithm-driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination?* (2020), <https://cdt.org/wp-content/uploads/2020/12/Full-Text-Algorithm-driven-Hiring-Tools-Innovative-Recruitment-or-Expedited-Disability-Discrimination.pdf> [https://perma.cc/9A74-WCCF].

⁶ Letter from ACLU et al. on Addressing Technology’s Role in Financial Services Discrimination to Consumer Fin. Prot. Bureau et al. (July 13, 2021), <https://www.aclu.org/letter/2020-07-13-coalition-memo-technology-and-financial-services-discrimination> [https://perma.cc/58YA-4YSQ]; Letter from Nat’l Fair Hous. All. on Request for Information and Comment on Financial Institutions’ Use of Artificial Intelligence, including Machine Learning to Fed. Rsrv. Sys. et al. (July 1, 2021), <https://nationalfairhousing.org/wp-content/uploads/2021/07/Federal-Banking-Regulator-RFI-re-AI-Advocate-Letter-FINAL-2021-07-01.pdf> [https://perma.cc/H529-V84L].

⁷ Ctr. for Democracy & Tech., *Algorithmic Systems in Education: Incorporating Equity and Fairness When Using Student Data* (Aug. 12, 2019), <https://cdt.org/insights/algorithmic-systems-in-education-incorporating-equity-and-fairness-when-using-student-data/> [https://perma.cc/PKY6-K6DR].

⁸ See Crystal Grant & Kath Xu, *Public Trust in Artificial Intelligence Starts with Institutional Reform*, ACLU (Sept. 17, 2021), <https://www.aclu.org/news/national-security/public-trust-in-artificial-intelligence-starts-with-institutional-reform> [https://perma.cc/MJP4-6VKC].

⁹ Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource, 86 Fed. Reg. 39081 (July 23, 2021).

that the NAIRR Task Force itself lacks sufficient expertise in the array of bias, equity, transparency, accountability, and civil rights/liberties issues that can flow from AI.¹⁰ Indeed, given the current structure of the Task Force and plan for its work, we have grave concerns that these harms are unlikely to be sufficiently addressed in the research, development, execution, and administration of the NAIRR.

First, while the members of the NAIRR Task Force and the panelists invited to appear before it to date are distinguished and qualified in many respects, they lack depth and range of expertise with civil rights/ liberties issues. This is not for a lack of such expertise in the field of AI. Indeed, there are many technical experts who have deeply engaged with how AI interplays with systemic discrimination on the basis of gender, race, disability or other protected characteristics, and yet such expertise is not sufficiently represented among the members of the Task Force.¹¹ Without including experts in bias, equity, transparency, accountability, and civil rights/liberties on the Task Force, these considerations are likely to be sidelined and ultimately viewed as an impediment to the development of AI. Furthermore, absent specific expertise on potential harms, an accurate and holistic “assessment of, and recommend[ed] solutions to, barriers to the dissemination and use of high-quality government data sets” is not possible.¹²

Second, the structure of and plan for the working groups is insufficient to adequately identify potential civil rights/liberties harms and create a plan to concretely address them in the development, execution, governance, and administration of the NAIRR. The Task Force’s work plan for considering these harms appears limited to the work of a separate “privacy, civil rights, and civil liberties” working group that only convenes in the final stage of the NAIRR Task Force’s assessment phase.¹³ Structural bias in big data and the harms it can cause to humans, particularly vulnerable or marginalized communities, cannot be an afterthought. Instead, the Task Force should embed consideration of bias, equity, transparency, accountability, and civil rights/liberties in its “goals and evaluations metrics” as well as in each of the working groups. In particular, the working groups on “governance models,” “data resources,” “user interfaces,” “educational tools,” and

¹⁰ See Letter from ACLU et al. to White House Off. of Sci. & Tech. Pol’y (July 13, 2021),

https://www.aclu.org/sites/default/files/field_document/2021-07-13_letter_to_white_house_ostp_on_centering_civil_rights_in_ai_policy_1.pdf [https://perma.cc/B3F8-Q29Q].

¹¹ Margaret Mitchell et al., *Diversity and Inclusion Metrics in Subset Selection*, 2020 Proc. AAAI/ACM Conf. on AI, Ethics, and Soc’y (Feb. 7–8, 2020), <https://dl.acm.org/doi/pdf/10.1145/3375627.3375832> [https://perma.cc/3J5W-TJDT]; Margaret Mitchell et al., *Model Cards for Model Reporting*, 2019 Conf. on Fairness, Accountability & Transparency (Jan. 29–31, 2019), https://arxiv.org/pdf/1810.03993.pdf?source=post_page [https://perma.cc/539H-S4QW]; Inioluwa Deborah Raji et al., *Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing*, 2020 Conf. on Fairness, Accountability & Transparency (Jan. 27–30, 2020), <https://dl.acm.org/doi/pdf/10.1145/3351095.3372873> [https://perma.cc/W43W-BHNU]; Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 2018 Conf. on Fairness, Accountability & Transparency (Feb. 23–24, 2018), <https://www.media.mit.edu/publications/gender-shades-intersectional-accuracy-disparities-in-commercial-gender-classification/> [https://perma.cc/343P-NRV9]; Deborah I. Raji & Jingying Yang, *ABOUT ML: Annotation and Benchmarking on Understanding and Transparency of Machine Learning Lifecycles*, 33rd Conf. on Neural Info. Processing Sys. (2019), <https://arxiv.org/pdf/1912.06166.pdf> [https://perma.cc/K5EQ-DLKM]; Timnit Gebru et al., *Datasheets for Datasets* (Mar. 19, 2020) (preprint), <https://arxiv.org/abs/1803.09010v7> [https://perma.cc/NL72-R75K].

¹² S. 3890, 116th Cong., 2d Sess. (2020), <https://www.congress.gov/116/bills/s3890/BILLS-116s3890is.xml> [https://perma.cc/3ZQJ-48X5].

¹³ Presentation, Nat. AI Rsch. Res. Task Force (Aug. 30, 2021), <https://www.ai.gov/wp-content/uploads/2021/09/NAIRR-TF-Presentations-08302021.pdf> [https://perma.cc/3WN3-JJBD].

“testing resources” must include robust consideration of these harms and ways to identify and address them. Preferably, each working group would include at least one member with substantive expertise in these harms so that their consideration can be incorporated in all the working groups’ recommendations.

Third, nothing in the Task Force’s publicly available agendas, minutes, and presentations suggests that there is a concrete plan for meaningful engagement with individuals and communities who have either experienced harm in the past from collection of data and deployment of AI, or who might suffer AI-related harms in the future. Although the Task Force appears to have heard from some experts on the topic of equity in access to AI education and funding, the Task Force also needs to include robust discussions about the potential harms that the NAIRR may pose to individuals and communities from whom data is collected and upon whom AI is deployed. End-users and impacted communities are often distinct and, unfortunately, the latter group is frequently left out of discussions about the use of data and AI. Each working group should create plans for how they will meaningfully solicit input on these topics from a variety of impacted communities and other experts on an ongoing basis. Presentation at a single meeting or brief consultations with experts will not provide for the deep considerations of bias, equity, transparency, accountability, and civil rights/liberties issues that is required to understand potential negative impacts on people and communities. To foster “public trust” and “prioritize areas of AI that solve problems for the public good”—issues that the panel on “Defining the Value Proposition and Intended Outcomes of a NAIRR” flagged as important¹⁴—it is critical that representatives of communities that may be impacted by the NAIRR, as well as other experts in the aforementioned harms, take part in the governance and decision-making about the NAIRR.

Several people and institutions possess both the type of technical expertise contemplated by the National AI Initiative Act, as well as expertise in the racial, gender, disability, and other biases that can arise in the collection and use of data and in development and deployment of AI. The Task Force could consider including civil rights/liberties organizations with particular expertise in data and AI, such as those listed below, in its working groups, list of consultants, and as candidates for participation in the governance structure of the NAIRR. Suggested organizations include but are not limited to: AI Now (<https://ainowinstitute.org/>), Algorithmic Justice League (<https://www.ajl.org/>), American Civil Liberties Union (<https://aclu.org/>), Artificial Intelligence, Policy, and Practice Initiative at Cornell University (<https://aipp.cis.cornell.edu/>), Center for Democracy and Technology (<https://cdt.org/>), Data 4 Black Lives (<https://d4bl.org/>), Georgetown Law School’s Center for Privacy and Technology (<https://www.law.georgetown.edu/privacy-technology-center/>), the Lawyers’ Committee for Civil Rights Under Law (<https://www.lawyerscommittee.org/>), the Leadership Conference (<https://civilrights.org/>), and Upturn (<https://upturn.org/>).

The Task Force should also solicit input from federal agencies that are tasked with working on bias, equity, and civil rights/liberties, including: the Civil Rights Division of the Department of Justice, Consumer Federal Protection Bureau, Department of Housing and Urban Development, Department of Labor, Equal Employment Opportunity Commission, Federal Trade Commission, and Office of Federal Contract Compliance Programs.

¹⁴ Meeting Summary, Nat’l AI Rsch. Res. Task Force (Aug. 30, 2021), <https://www.ai.gov/wp-content/uploads/2021/09/NAIRR-TF-Meeting-Minutes-08302021.pdf> [https://perma.cc/AE2W-9CUQ].

III. Task Force Emphasis on Large-Scale Computing Resources Reifies a Type of AI That Is Especially Costly for Society

We believe that the Task Force has, to date, over-emphasized the role of large-scale compute in its considerations. This leads to the myopic view that building bigger and more expansive compute is necessary for positive advances in the AI field. The Task Force needs to consider recent scholarship questioning the tradeoffs associated with building larger models as it decides where to invest limited government resources.¹⁵ The financial¹⁶ and environmental¹⁷ costs for small gains in accuracy are substantial in the places where increasing computational power works—and in many domains (e.g., tabular data), adding compute or ever more complex models may not do much to improve performance.¹⁸ The hoovering up of ever-larger datasets turns them into “stochastic parrots” of some of the most toxic information within that data¹⁹ and can lead to increasingly exploitative data collection practices.²⁰ The private sector²¹ and the government itself²² have now invested many billions in capital in ever more expansive AI, yet many of the most vexing problems remain, and little of this investment addresses the harms these models may inflict on society. The Task Force should consider whether investing in the compute-race is the best use of its limited resources. Some of the biggest proponents of the compute-heavy AI that the Task Force envisions are now questioning the emphasis on these models,²³ and the Task Force should as well.

¹⁵ Emily M. Bender et al., *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, Proc. of the 2021 ACM Conf. on Fairness, Accountability, and Transparency 610 (2021), <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922> [https://perma.cc/MD4V-EYAN] (critiquing large language models); Song Han, Huizi Mao, & William J. Dally, *Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding*, Int’l Conf. on Learning Representations (2016), <https://arxiv.org/abs/1510.00149> [https://perma.cc/8YL9-JAHN] (focusing on high efficiency as opposed to ever-larger models).

¹⁶ See *The Staggering Cost of Training SOTA AI Models*, Synced (June 27, 2019), <https://medium.com/syncedreview/the-staggering-cost-of-training-sota-ai-models-e329e80fa82> [https://perma.cc/3D9Q-G77F].

¹⁷ See Mark Labbe & Ronald Schmelzer, *AI and Climate Change: The Mixed Impact of Machine Learning*, EnterpriseAI (Aug. 31, 2021), <https://searchenterpriseai.techtarget.com/feature/AI-and-climate-change-The-mixed-impact-of-machine-learning> [https://perma.cc/54AE-MKC7]; Karen Hao, *Training a Single AI Model Can Emit as Much Carbon as Five Cars in Their Lifetimes*, MIT Tech. Rev. (June 6, 2019), <https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes> [https://perma.cc/CX6R-5C5P].

¹⁸ Arlind Kadra et al., *Regularization is All You Need: Simple Neural Nets Can Excel on Tabular Data* (2021), <https://arxiv.org/pdf/2106.11189.pdf> [https://perma.cc/59WH-8BX2] (preprint).

¹⁹ Bender et al., *supra* note 15.

²⁰ See Eric Null, Iseuda Oribhabor, & Willmary Escoto, *Data Minimization: Key to Protecting Privacy and Reducing Harm*, Access Now (2021), <https://www.accessnow.org/cms/assets/uploads/2021/05/Data-Minimization-Report.pdf> [https://perma.cc/99GU-J25C]; Phil Jones, *Refugees Help Power Machine Learning Advances at Microsoft, Facebook, and Amazon*, Rest of World (Sept. 22, 2021), <https://restofworld.org/2021/refugees-machine-learning-big-tech> [https://perma.cc/3MKF-J46L].

²¹ See Karen Hao, *Inside the Fight to Reclaim AI from Big Tech’s Control*, MIT Tech. Rev. (June 14, 2021), <https://www.technologyreview.com/2021/06/14/1026148/ai-big-tech-timnit-gebru-paper-ethics> [https://perma.cc/X4ZG-RZ4B].

²² See Jon Harper, *Federal AI Spending to Top \$6 Billion*, Nat’l Def. (Feb. 10, 2021), [https://www.nationaldefensemagazine.org/articles/2021/2/10/federal-ai-spending-to-top-\\$6-billion](https://www.nationaldefensemagazine.org/articles/2021/2/10/federal-ai-spending-to-top-$6-billion) [https://perma.cc/Q24J-3H8D].

²³ See Gary Marcus, *Deep Learning: A Critical Appraisal* (2018), <https://arxiv.org/pdf/1801.00631.pdf> [https://perma.cc/Q5D8-WR86]; François Chollet (@fchollet), Twitter (Dec. 18, 2017), <https://twitter.com/fchollet/status/942733414788190209> [https://perma.cc/6E65-VJGN] (“For most problems where

For the same reasons, we believe the establishment of the NAIRR should focus on offering an alternative to the data- and compute-hungry applications that are the focus of many industry and research labs—most of whom already have substantial resources at their disposal and need no additional government support. Instead, the Task Force should emphasize shared data access to researchers, journalists, and advocates who are frequently shut out from access to datasets held by companies or public sector entities. The scarce resource for underrepresented researchers is not technology infrastructure—it is access to relevant, real-world data that companies and public agencies guard closely. Even for those focused on the use of deep learning, there is a longstanding problem with the data sets available to researchers,²⁴ especially those who are not associated with large technology companies and who may be able to shed light into the risks of this technology or develop new innovations.

We will spend the remainder of this paper focusing on the NAIRR as a shared data repository. We ask that the Task Force especially consider data that would most help civil rights and society groups to inspect and understand widely used AI, the privacy and ethical implications of data collection and/or acquisition, and the risks associated with the NAIRR.

IV. Shared Data that Would Help Advance Civil Rights Concerns and the Ability to Inspect

To advance the protection of civil rights, any shared data repository should contain a number of mandatory datasets as well as accompanying requirements for these datasets. We recommend that the NAIRR Task Force require the inclusion of datasets used by local, state, or federal agencies that have elected to build predictive models/automated decision systems.

Predictive models by these government actors have a wide-ranging impact on individual lives. Examples of predictive models currently being used by local government agencies include pretrial risk assessment tools and child welfare screening tools. Pretrial risk assessment tools are often built on datasets that include demographic information and information about an individual's criminal legal system history.²⁵ Child welfare screening tools can draw from datasets that include demographic information and child welfare, jail, juvenile probation, behavioral health, and birth records.²⁶ Journalists and civil rights groups have already raised equity concerns about both.²⁷

deep learning has enabled transformationally better solutions (vision, speech), we've entered diminishing returns territory in 2016-2017.”)

²⁴ See Amandalynne Paullada et al., *Data and Its (Dis)contents: A Survey of Dataset Development and Use in Machine Learning Research*, NeurIPS 2020 Workshop, <https://arxiv.org/pdf/2012.05345.pdf> [https://perma.cc/7B6G-8CLW].

²⁵ Mapping Pretrial Injustice, *Inputs: Variables*, <https://pretrialrisk.com/the-basics/pretrial-risk-assessment-instruments-prai/inputs-variables> [https://perma.cc/F58Z-5YLB].

²⁶ See Oregon Dep't of Hum. Servs., Oregon DHS Safety at Screening Tool—Development and Execution at 4 (2019), <https://www.oregon.gov/DHS/ORRAI/Documents/Safety%20at%20Screening%20Tool%20Development%20and%20Execution%20Report.pdf> [https://perma.cc/X4ZT-YR56]; Rhema Vaithianathan et al., Allegheny Family Screening Tool: Methodology, Version 2, at 3 (2019), <https://www.alleghenycountyanalytics.us/wp-content/uploads/2019/05/Methodology-V2-from-16-ACDHS-26 PredictiveRisk Package 050119 FINAL-7.pdf> [https://perma.cc/KLJ9-AHAB].

²⁷ Julia Angwin et al., *Machine Bias*, ProPublica (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [https://perma.cc/ESW2-B9YQ]; Anjana Samant et al., *Family Surveillance by Algorithm: The Rapidly Spreading Tools Few Have Heard Of*, ACLU (Sept. 29, 2021), <https://www.aclu.org/news/womens-rights/family-surveillance-by-algorithm-the-rapidly-spreading-tools-few-have-heard-of> [https://perma.cc/M55Z-F2XL].

One aspirational goal for NAIRR would be to offer resources for community-based journalistic and academic researchers to provide oversight on powerful private sector actors that use AI over exclusive/proprietary data. These institutions often shut out marginalized or underrepresented groups looking to understand their impact. A recent example is Facebook’s decision to block “good-faith research in the public interest” and shut down the accounts of New York University researchers who were studying how the platform’s political advertisers target different users.²⁸ For these reasons, the government should encourage private actors to offer datasets that would help researchers evaluate the fairness risks of tools that rely on these datasets. The NAIRR Task Force should consider how to provide incentives and support to independent researchers who are engaged in the oversight of privatized AI/ML systems that have significant social impact.

To ensure the usefulness of datasets, the NAIRR Task Force should set standards for data documentation and require civil rights assessments of both the use case and data provenance, as described further in the sections below.²⁹ To facilitate audits, the datasets should include protected class data such as race and gender when legally permitted, and when it is not, sufficient data to infer protected class status using standardized methods.³⁰ Contributors of datasets should be required to document inherent flaws of their data (e.g., selection bias) and discuss the possible creation of datasets that overcome these flaws. Datasets should also include qualitative data, including feedback from listening sessions with impacted communities and from tool users such as agency workers.³¹ Finally, if the NAIRR Task Force determines that the submitted data or methods are too flawed, it should be prepared to ask an entity to halt or refrain from deploying the predictive model built on that data and should issue an explanatory statement recommending the technology not be used. We hope the creation of NAIRR, if open to a broader set of researchers, can encourage more impact evaluations like the National Institute of Standards and Technology’s Face Recognition Vendor Test report or the research reports resulting from the Criminal Justice Administrative Records System.³²

²⁸ See Letter from Samuel Levine, Acting Dir. of the Bureau of Consumer Prot., to Facebook (Aug. 5, 2021), <https://www.ftc.gov/news-events/blogs/consumer-blog/2021/08/letter-acting-director-bureau-consumer-protection-samuel> [https://perma.cc/68PD-HA4N]; Ethan Zuckerman, *Facebook Cares About Privacy—But Only If You’re an Advertiser*, *The Atlantic* (Aug. 6, 2021), <https://www.theatlantic.com/technology/archive/2021/08/facebook-only-cares-about-privacy-advertisers/619691/> [https://perma.cc/6RWH-YBY4].

²⁹ See Timnit Gebru et al., *Datasheets for Datasets* (Mar. 19, 2020) (preprint), <https://arxiv.org/abs/1803.09010v7> [https://perma.cc/NL72-R75K]; *The Dataset Nutrition Label*, Data Nutrition Project, <https://datanutrition.org/labels/> [https://perma.cc/GD9C-LHXQ]; Bao et al., *supra* note 3.

³⁰ Regulation B of the Equal Credit Opportunity Act, for example, does not allow the use or collection of protected characteristics, such as race or gender, by financial institutions except for mortgage decisions. See, e.g., Equal Credit Opportunity Act (Regulation B) Ethnicity and Race Information Collection, 12 C.F.R. § 1002 (2017), <https://www.federalregister.gov/d/2017-20417> [https://perma.cc/VR8C-D2AH]. The standard approach in the credit and lending industry is to use publicly available information as a proxy for unidentified race and ethnicity. Consumer Fin. Prot. Bureau, *Using Publicly Available Information to Proxy for Unidentified Race and Ethnicity* (2014), https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf [https://perma.cc/HQ2E-PV63].

³¹ See Kelley Fong, *Getting Eyes in the Home: Child Protective Services Investigations and State Surveillance of Family Life*, 85 *Am. Socio. Rev.* 610 (2020), <https://journals.sagepub.com/doi/10.1177/0003122420938460> [https://perma.cc/P5V7-DJDA]; Virginia Eubanks, *A Child Abuse Prediction Model Fails Poor Families*, *Wired* (Jan. 15, 2018), <https://www.wired.com/story/excerpt-from-automating-inequality/> [https://perma.cc/786P-JAAM].

³² Patrick Grother, Mei Ngan & Kayee Hanaoka, Nat’l Inst. of Standards & Tech., NISTIR 8280, *Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects* (Dec. 2019), <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf> [https://perma.cc/PTQ8-UNVE]; Research Reports,

V. Civil Rights Impact Assessment Requirements for Data Receipt and Use

We encourage the NAIRR Task Force to ensure that those who receive data from the program are held to a high standard. This may include the use of civil rights impact assessments and other mechanisms of assurance in any implementation of the NAIRR before providing data, but also by those who intend to use the NAIRR’s datasets or research with it. These mechanisms must meaningfully include input from impacted people and communities in their design and execution.³³ As we have written elsewhere, we believe such impact assessments should be conducted according to standards that set out necessary evaluation points, and at a minimum should require: regular evaluation for discriminatory effects throughout the conception and development of any model based on the data, and—if not terminated during development due to unacceptable impacts or for other reasons—in its implementation and use; proactive searches for and adoption of less discriminatory alternatives; continuing assessments of whether the data used in training technologies is representative and accurate; and that the technologies measure lawful and meaningful attributes and seek to predict valid target outcomes.³⁴

While some entities are voluntarily evaluating their AI systems, there is too little transparency in the documentation and publication of the resulting impact assessments. The Task Force could help solve this issue by enhancing researcher access to previously unevaluated AI datasets, and it should also work to ensure that any uses of the data come with documentation, transparency, and accountability requirements. When an impact assessment is conducted, it is important that the Task Force require information about the evaluation be made publicly available, including information about the content and reasoning behind the evaluation, who is conducting the evaluation, and what their relationship is to the entity being evaluated, if the evaluation is not conducted by a government agency itself.

If educational materials will be provided to those who use or access the NAIRR, those materials should go beyond mechanical and technical concepts; they should also address issues inequity, AI ethics, and the potential disparate impact of AI.³⁵ Educational resources should also include details on how to conduct audits for fairness and disparate impact of AI tools.

Crim. Just. Admin. Rec. Sys., <https://cjars.isr.umich.edu/overview/research/research-reports/> [https://perma.cc/Z2HD-QCSW].

³³ See Dillon Reisman et al., AI Now Inst., *Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability* 18–20 (Apr. 18, 2018), <https://ainowinstitute.org/aiareport2018.pdf> [https://perma.cc/JQY8-VB4L] (“[Entities] should ensure that affected communities are able to suggest researchers that they feel represent their interests, and should work with researchers to ensure that these communities have a voice in formulating the questions that are asked and addressed by research and auditing.”).

³⁴ See *ACLU Comment on NIST’s Proposal for Managing Bias in AI*, ACLU (Sept. 10, 2021), <https://www.aclu.org/letter/aclu-comment-nists-proposal-managing-bias-ai> [https://perma.cc/B58E-VXF5].

³⁵ Maria Kasinidou et al., *Educating Computer Science Students About Algorithmic Fairness, Accountability, Transparency and Ethics*, 1 Proc. of the 26th ACM Conf. on Innovation & Tech. in Comp. Sci. Educ. 484 (2021), <https://dl.acm.org/doi/abs/10.1145/3430665.3456311> [https://perma.cc/H53F-KBR2]; Veronika Bogina et al., *Educating Software and AI Stakeholders About Algorithmic Fairness, Accountability, Transparency, and Ethics*, Int’l J. of Artificial Intel. in Educ. 1 (2021), <https://link.springer.com/article/10.1007/s40593-021-00248-0> [https://perma.cc/RZ5U-3A9K]; Anna Lauren Hoffman & Katherine Alejandra Cross, *Teaching Data Ethics: Foundations and Possibilities from Engineering and Computer Science Ethics Education*, <https://digital.lib.washington.edu/researchworks/bitstream/handle/1773/46921/TeachingDataEthicsFoundations-Hoffmann-Cross.pdf> [https://perma.cc/HH92-7YPF].

VI. Privacy Implications of Data Collection and/or Acquisition

The NAIRR aims to make creating AI tools more accessible to researchers and other AI practitioners, and this may include making data available for models to be trained and tested. This raises questions of what kinds of data the Task Force believes should be made available. The Task Force will need to make clear the kinds of data it believes are appropriate to be made widely available and what AI tools are appropriate to be developed based on this data. For example, the Task Force will need to decide whether images of faces will be among data it makes available. While numerous studies have detailed that current facial recognition technology is in need of improvement in accuracy, some organizations, including the ACLU, are opposed to the creation of these tools because of their potential use for pervasive surveillance even after technological biases have been addressed. Similarly, the inclusion of public data on arrests or crime rates may be controversial because of the known history of racial and ethnic bias within this data. The Task Force must take care to ensure that data selected to be included in this database cannot be used to perpetuate harm and increase surveillance.

VII. Datasets Present Risks

Distributing access to any dataset may present risks to individuals whose data are present, or to communities of people similar to those present in the dataset. The Task Force should consider these risks to data subjects and have clear plans for how to mitigate them as it looks to expand access to datasets. Examples of risks include identity theft, harassment, discrimination, unwarranted legal scrutiny, surveillance, and financial harm data harms.³⁶ Such harms can arise from release of medical information, location history, personal imagery, financial records, employment details, sexual history, religious information, and more. Any large, socially relevant dataset can contain sensitive information that can cause harm: either to people or communities represented in the datasets, or to people or communities subject to systems trained on the dataset. The NAIRR will likely collect, hold, and redistribute datasets, and thus has responsibilities as a data steward.

VIII. Redress for Harm and Responsibility for Researchers

It can be difficult for an individual or group to know that a harm they have experienced is due to a dataset released for research. It is also unclear how to identify parties who have been harmed, or how to meaningfully redress harms when they are identified. When a researcher causes harm through misuse of a dataset, inadvertently or on purpose, it is not clear what consequences or liability they face, or how it would be enforced. The Task Force should address these concerns and develop policies and practices to minimize harm, provide redress when it occurs, and structure liability for researchers.

IX. Examples of Attempts at Responsible Data Management

Managers of other large, research-oriented datasets have tried to tackle similarly significant challenges related to privacy and other harms. An example is the National Center for Health Statistics Research Data Center's documentation of anti-disclosure rules for researchers,³⁷ which include both institutional and individual consequences, and association with legal liability regimes

³⁶ Danielle Keats Citron & Daniel J. Solove, *Privacy Harms*, GWU Legal Stud. Rsch. Paper No. 2021-11 (Feb. 9, 2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3782222 [https://perma.cc/N2H3-6KYH].

³⁷ Nat'l Ctr. for Health Stat. Rsch. Data Ctr., *Preventing Disclosure: Rules for Researchers* (2020), <https://www.cdc.gov/rdc/data/b4/Disclosure-Manual-v2.5.pdf> [https://perma.cc/WEM3-3SVW].

like the Confidential Information Protection and Statistical Efficiency Act.³⁸ Another example is the U.S. Census Bureau's responsibility for distributing large amounts of analyzable data with potentially serious side effects for the people involved. Recent attempts at safer data distribution policies include differential privacy and other statistical safeguards.³⁹ These privacy preserving techniques may also reduce the accuracy of the data or may limit its applicability for certain kinds of use. In addition, Duke University has published a dataset known as DukeMTMC aimed at improving vision analysis, person-recognition, and object tracking. Duke later took down that dataset⁴⁰ in response to a report about its widespread use⁴¹ and associated harms, but the data continues to be used despite the retraction.⁴² While these examples include attempts to mitigate harm, the efficacy of these mitigations is unclear, as is researcher liability when datasets are misused. The Task Force should ensure that the NAIRR does better at protecting data subjects, redressing harms, and ensuring consequences for researchers whose work causes harm.

We thank you for considering our suggestions.

Sincerely,

The American Civil Liberties Union

³⁸ Confidential Information Protection and Statistical Efficiency Act of 2002, 44 U.S.C. §§ 3501-3521, <https://www.eia.gov/cipsea/cipsea.pdf> [https://perma.cc/G83Z-JZQQ].

³⁹ Simson L. Garfinkel et al., U.S. Census Bureau, Differential Privacy at the US Census Bureau: Status Report (Jan. 27, 2020), <https://csrc.nist.gov/CSRC/media/Projects/pec/documents/stppa-01-20200127-talk03-Garfinkel-diff-priv-census.pdf> [https://perma.cc/AG7K-JX5R]; U.S. Census Bureau, Statistical Safeguards, https://www.census.gov/about/policies/privacy/statistical_safeguards.html [https://perma.cc/TA4P-QHCP] (last visited Sept. 30, 2021).

⁴⁰ Jake Satsky, *A Duke study recorded thousands of students' faces. Now they're being used all over the world*, Duke Chron. (June 11, 2019), <https://www.dukechronicle.com/article/2019/06/duke-university-facial-recognition-data-set-study-surveillance-video-students-china-uyghur> [https://perma.cc/X2L8-V54L].

⁴¹ *Duke MTMC Dataset*, Exposing.ai, https://exposing.ai/duke_mtmc/ [https://perma.cc/VEP7-D7J8] (last visited Sept. 30, 2021).

⁴² Kenny Peng, *Facial recognition datasets are being widely used despite being taken down due to ethical concerns. Here's how.*, Freedom to Tinker (Oct. 21, 2020), <https://freedom-to-tinker.com/2020/10/21/facial-recognition-datasets-are-being-widely-used-despite-being-taken-down-due-to-ethical-concerns-heres-how/> [https://perma.cc/FD93-LNQU].

Federal Register Notice 86 FR 46278, <https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

American Psychological Association (APA)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1, 2021

[REDACTED]
[REDACTED]
[REDACTED]

Submitted electronically [REDACTED]

RE: 86 FR 39081: Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource

Dear Ms. Wigen:

The American Psychological Association (APA) is grateful for the opportunity to respond to the White House Office of Science and Technology Policy (OSTP) and the National Science Foundation's (NSF) [Request for Information \(RFI\): Implementation Plan for a National Artificial Intelligence Research Resource \(86 FR 39081\)](#). We thank OSTP and NSF for seeking wide input to inform federal activities attempting to democratize advanced cyberinfrastructure in the US.

APA is the leading scientific and professional organization representing psychology in the US, with more than 122,000 researchers, educators, clinicians, consultants, and students as its members and affiliates. APA's mission is to promote the advancement, communication, and application of psychological science and knowledge to benefit society and improve lives. As a broad field, psychology can contribute to the development of artificial intelligence (AI), enhance its positive impacts for individuals and society, and reduce unintended negative consequences. Our association is also committed to increasing access to scientific and educational tools for underrepresented research institutions and investigators, and to the ethical and responsible development of AI systems. We applaud the National AI Research Resource (NAIRR) Task Force's efforts to develop representative datasets, appropriate educational tools, and support mechanisms to improve equitable engagement with AI resources in the US.

The following are responses to the questions prompted by the implementation roadmap:

Question 2: Which capabilities and services provided through the NAIRR should be prioritized?

The Task Force has proposed several potential capabilities and services necessary to maintain a shared computing infrastructure and facilitate equitable access to resources for researchers across the country. Among these capabilities, it is essential to prioritize the appropriate choice and use of metadata – this will enhance the understanding, organization, and use of curated

datasets under NAIRR.¹ However, sufficiently appropriate metadata is only one necessary ingredient to meet the work plan for NAIRR as it is enumerated. For example, Section E references “high-quality government data sets” where “high-quality” implies that there is appropriate measurement, in place of mainstream measures that may be superficially appealing. The [AERA/APA/NCME Standards for Educational and Psychological Testing and Society for Industrial and Organizational Psychology \(SIOP\) Principles for the Validation and Use of Personnel Selection Procedures](#) are robust examples of valuable contributions from psychological science, including decades of psychometric and statistical research relevant to the successful execution of NAIRR’s workplan and priorities.

Per Sections D-G, human factors (HF) psychologists² are critical to the incorporation of usability, and human-systems integration into AI and data-sharing technologies.³ HF researchers are necessary to ensure the AI data-sharing systems⁴ and cyberinfrastructures developed will be useful, accessible, privacy-assuring, and fair to all parties involved. Experts who have developed, revised, implemented, and maintained taxonomies, including psychologists (and the rich ontologies supporting them) must be included to construct dataset platforms that are strong, scalable, useful, and proven capable of preserving privacy.⁵

Question 3: How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Psychological science is vital to the ethical and responsible development of AI. Some AI systems have been trained on large data sets of human attributes or demographics that integrate biases related to gender, race, and other characteristics. These systems then replicate the biases in their interactions with humans, with implications for equity and fairness. Psychologists’ research on the various forms of resulting bias and the detrimental impacts are being used to develop data sets that are less biased and AI systems that can detect and compensate for biases in data.

¹ Blask, K., Gerhards, L., & Jalynskij, M. (2020, February 1). *Metadata in Psychology 2.0: What researchers really need – Study description of the data referring to the online survey conducted in the BMBF-funded project PsyCuraDat*. PsychArchives. <http://dx.doi.org/10.23668/psycharchives.2757>

² Brée D.S. (1988). *Artificial Intelligence and Cognitive Psychology: A New Look at Human Factors*. In: van der Veer G.C., Mulder G. (eds) *Human-Computer Interaction*. Springer. https://doi.org/10.1007/978-3-642-73402-1_17

³ Tenopir, C., Rice, N.M., Allard, S., Baird, L., Borycz, J., Christian, L., Grant, B., Olendorf, R., Sandusky, R.J., (2020). Data sharing, management, use, and reuse: Practices and perceptions of scientists worldwide. *PLOS ONE*, 15(3). <https://doi.org/10.1371/journal.pone.0229003>

⁴ Tsuji, S., Bergmann, C., & Cristia, A. (2014). Community-Augmented Meta-Analyses: Toward Cumulative Data Assessment. *Perspectives on Psychological Science*, 9(6), 661–665. <https://doi.org/10.1177/1745691614552498>

⁵ Hesse, B. W. (2018). Can psychology walk the walk of open science? *American Psychologist*, 73(2), 126–137. <https://doi.org/10.1037/amp0000197>

Given evidence that AI can reproduce discrimination and bias against individuals and groups, it is imperative to leverage psychological science and examine people's expectations about and reactions to the fairness and potential discrimination of AI versus human agents. An emerging line of research suggests that people expect AI to be less biased than humans in some cases, and are less outraged when they learn of bias from an AI versus human actors.⁶ Algorithms appear less discriminatory than humans, perhaps incorrectly engendering trust and comfort from human users. Given the massive and increasing influence of AI on people's lives, it is critical to better appreciate how people understand and react to such influence, especially when the AI is perceived to be biased or unfair.

There are some fundamental research opportunities the AI research community must investigate. AI Ethics and Psychology is an evolving discipline essential to the study of how AI learns from society and humans⁷ and how AI makes consequential decisions in critical settings.⁸ Studies have demonstrated that AI automatically learns implicit biases from language corpora and accordingly perceives the world in a biased manner. These implicit biases that have been documented in social psychology for decades include racial, gender, sexuality, ability, and age attitudes.⁹¹⁰ Moreover, these findings provide insights about how language might be impacting the social cognition of both AI and humans.

There are, additionally, ethical implications for what AI learns, how AI learns, and AI's subsequent decision-making. For example, developing transparency enhancing algorithms for measuring and simulating AI bias and equity would make it possible to analyze the ethical implications of AI in a variety of domains including natural language and computer vision.¹¹ Alternatively, these AI methods could examine and analyze current and historical social and human cognition.¹² This research program would allow for understanding how AI is co-evolving

⁶ Jago, A. S., & Laurin, K. (2021). Assumptions About Algorithms' Capacity for Discrimination. *Personality and Social Psychology Bulletin*. <https://doi.org/10.1177/01461672211016187>

⁷ Caliskan, A., Bryson, J.J., & Narayanan, A., (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183-186. [10.1126/science.aal4230](https://doi.org/10.1126/science.aal4230)

⁸ Pandey, A., & Caliskan, A., (2021). *Disparate Impact of Artificial Intelligence Bias in Ridehailing Economy's Price Discrimination Algorithms*. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society. 822-833.

⁹ Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4–27. <https://doi.org/10.1037/0033-295X.102.1.4>

¹⁰ Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74(6), 1464–1480. <https://doi.org/10.1037/0022-3514.74.6.1464>

¹¹ Steed, R., & Caliskan, A. (2021). A set of distinct facial traits learned by machines is not predictive of appearance bias in the wild. *AI Ethics* 1, 249–260. <https://doi.org/10.1007/s43681-020-00035-y>

¹² Caliskan, A., & Lewis, M. (2020, July 16). Social biases in word embeddings and their relation to human cognition. <https://doi.org/10.31234/osf.io/d84kg>

with humanity, as AI is shaping society and impacting individuals' lives in an accelerated manner and at an unprecedented scale. While beyond the scope of this Request for Information, there remains no comprehensive regulation for auditing how AI impacts equity and fairness in democratic societies.¹³ Consequently, these promising research areas of computer and information science contribute to data-driven policy making and law while having implications for psychology, political science, sociology, linguistics, and philosophy.

Question 6: Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

NAIRR's ability to democratize access to AI R&D may rely upon the ethical and responsible development of AI systems. Ensuring ethical and responsible AI R&D will require experts who understand not only the meaning of equity, fairness, and bias, but also the nature of AI technologies and algorithms¹⁴ – and importantly, the samples and settings to which they apply.¹⁵ Computer scientists and engineers must enlist the social, behavioral, and psychological sciences to ensure the development less harmful AI systems that have the potential to spark positive social change.

APA looks forward to a close collaboration as this plan is further developed and implemented. Please contact with any questions, I can be reached at [REDACTED]

Sincerely,

[REDACTED]

Mitch Prinstein, PhD
Chief Science Officer

¹³ Caliskan, A. (2021, May 10). *Detecting and mitigating bias in natural language processing*. Brookings Institution. <https://www.brookings.edu/research/detecting-and-mitigating-bias-in-natural-language-processing/>

¹⁴ McLennan, S., Fiske, A., Celi, L.A. et al. 2020). An embedded ethics approach for AI development. *Nature Machine Intelligence* 2, 488–490. <https://doi.org/10.1038/s42256-020-0214-1>

¹⁵ Sloane, M., & Moss, E. (2019). AI's social sciences deficit. *Nature Machine Intelligence* 1, 330–331. <https://doi.org/10.1038/s42256-019-0084-6>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Anthropic

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

ANTHROPIC

September 30, 2021

Submitted by email

Re: Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource

Anthropic welcomes the opportunity to provide feedback to the Office of Science and Technology Policy (OSTP) and the National Science Foundation (NSF) in response to a Request for Information (RFI) on the development of the National Artificial Intelligence Research Resource (NAIRR). Anthropic is an AI safety and research company working to build reliable, interpretable, and steerable artificial intelligence (AI) systems. We're an organization with backgrounds in research, engineering, and policy, and we approach AI development from a cross-disciplinary perspective.

We believe progress in AI benefits from a broad range of stakeholders understanding the technology and its effects, and participating actively in its development. In addition to democratizing AI resources for the research and academic community, the NAIRR provides an opportunity for the U.S. government to measure and monitor progress in AI research and development (R&D), by virtue of creating a shared infrastructure that many stakeholders can use. With detailed information about the sorts of computationally-intensive projects researchers want to do, the U.S government will be better able to identify future research directions and funding needs¹.

We appreciate the opportunity to share our perspective on how the Task Force may wish to implement the technical and organizational design of the NAIRR. At a high level, we propose the NAIRR:

- Serve as a centralized marketplace for access to compute resources, made available by U.S. cloud providers;
- Prototype a system to access easily usable compute facilitated by supercomputers from the Department of Energy;

¹ Whittlestone, J., & Clark, J. (2021). Why and How Governments Should Monitor AI Development. *arXiv*. arxiv.org/pdf/2108.12427.pdf

- Provide funding for a variety of AI research projects through a two-track proposal review process: one for Principal Investigator(s) (PIs) from U.S. academic institutions, and another reserved for a handful of large-scale, computationally-intensive research projects;
- Support the diversification of AI research talent in the U.S. by allocating a portion of funds to PIs that have not previously received NAIRR funding for work in AI;
- Encourage a significant portion of research conducted in the NAIRR be contributed to open source platforms in the form of code and research published as preprints;
- Develop a set of success criteria for the NAIRR, in order to generate information about how useful it is, and provide a means to iterate on program and technical design choices as needed.

Summary

We must ensure that the academic community has sufficient access to the infrastructure and services required for research into AI, much of which has become computationally-intensive. To achieve this, we recommend that the Task Force consider a phased approach to the development of the NAIRR, prioritizing the most critical and readily-available components of a shared compute resource first, in order to enable researchers to be able to replicate experiments done at the frontier of resource-intensive AI technology development. As these foundational components are made available, the NAIRR can be enhanced in subsequent iterations with the addition of important, though less time-sensitive, services and capabilities.

If successful, the NAIRR and its accompanying programs will democratize access to AI R&D by making infrastructure technology available to a broader range of academic and research institutions.

Compute: While strengthening the AI ecosystem will require investments in a number of important areas (e.g. education and training, dataset availability, etc.), arguably none are more strategic or immediately pressing as access to compute resources. OpenAI notes that the amount of compute used to train some large AI systems has doubled every 3.4 months since 2012, compared to the previous two-year doubling period for AI systems². This indicates that access to large-scale computational resources has been critical in the creation of various AI breakthroughs — a point also made by Richard Sutton, an AI professor and pioneer of reinforcement learning, in his essay ‘The Bitter Lesson’³.

In this essay, Sutton argues that “the biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin.” He concludes that “one thing that should be learned from the bitter lesson is the great power of general purpose methods, of methods that continue to scale with increased computation even as the available computation becomes very great.” This suggests that many AI

² Amodei, D., & Hernandez, D. (2018, May 16). *AI and Compute*. OpenAI. <https://openai.com/blog/ai-and-compute/>

³ Sutton, R. (2019, March 13). *The Bitter Lesson*. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>

breakthroughs not only *benefit* from access to large-scale computational infrastructure, but the discovery of these breakthroughs may *require* access to increasingly large amounts of compute to be found. As this trend continues, the infrastructure required to develop these systems will become increasingly cost-prohibitive for all but a handful of entities, most of which are private corporations.

To ensure researchers can fully participate in AI R&D, the NAIRR must help close the gap in infrastructure access between academic and industrial research labs. For example, at Anthropic there are multiple GPUs per researcher, whereas when we talk to academic labs we tend to find the inverse: that there are sometimes multiple researchers per GPU. The Task Force may wish to conduct an analysis of compute requirements for large-scale AI projects in both academia and industry to quantify this gap and determine the amount of funding necessary to support AI R&D through the NAIRR. On the technical side, a prototype NAIRR could serve as a centralized marketplace to acquire compute resources, leveraging the best infrastructure from U.S. cloud providers and existing public sector resources such as supercomputers managed by the Department of Energy.

Research Funding: With funding made available through a two-track proposal review process, successful grantees could use funds in the form of marketplace credits to acquire the infrastructure services that best meet the needs of their particular research efforts. One track could look to existing review processes such as National Science Foundation grants⁴ or Broad Agency Announcements (BAAs) from the Defense Advanced Research Projects Agency (DARPA)⁵, and solicit proposals from PIs in a method that is rigorous, transparent, and competitive. A parallel track would be used to fund a handful of large-scale, computationally-intensive projects and would be subject to a more robust review process. To encourage an equitable distribution of funding, we recommend reserving a portion of funds each year for PIs that have not previously received NAIRR grants for AI research.

1. What options should the Task Force consider for any of roadmap elements A through I above, and why?

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success

Prior to the development of the NAIRR, the Task Force should establish a set of metrics to evaluate at predefined intervals during the lifecycle of the NAIRR to assess the success of the initiative. Some preliminary metrics the NAIRR may consider evaluating include: the scale of AI projects academia is able to develop as a result of resources unlocked by the NAIRR, the usability of NAIRR resources, the number of PIs receiving a research grant for the first time, the

⁴ National Science Foundation. *Merit Review*. https://www.nsf.gov/bfa/dias/policy/merit_review/

⁵ Defense Advanced Research Projects Agency. *Office-wide Broad Agency Announcements*. <https://www.darpa.mil/work-with-us/office-wide-broad-agency-announcements>

geographic diversity of PIs receiving funding, the number of publications enabled by the NAIRR, and the number of open source contributions to NAIRR-enabling infrastructure.

One metric of success would be to determine whether or not the NAIRR can help enable academic researchers to build models that match the resource scale (amount of compute used in training, size of datasets trained on, and so on) of those built by industry — this is important for developing AI research capacity in academia that is able to replicate (and therefore, critique and analyze) the systems being developed in industry.

Adoption and use of the NAIRR will depend heavily on its usability, which should be regularly measured as an additional indicator of success. The Task Force may implement various usability metrics indicating how long it takes researchers to get up and running with an experiment after receiving funding. Usage patterns and researcher surveys might also reflect the usability of the NAIRR, leading to technical or program design changes in later implementations. We elaborate further on how Anthropic enhances usability for our own infrastructure in Section 2, below.

The NAIRR presents a unique opportunity to help diversify the range of actors contributing to AI R&D. Providing funding and technical tools to researchers builds on a long partnership between the federal government and academia in discovering the scientific breakthroughs that power the U.S. economy. It is critically important that the NAIRR provide opportunities to newer entrants in the AI research community in order to both diversify the current talent pool and create a broader foundation for the U.S. research community.

To do this, we recommend the NAIRR reserve a portion of each year’s funding (e.g. 10% of total grant dollars) for PIs that have not previously received NAIRR grants for AI research. This means that, beginning in its second year of operation, the NAIRR would start setting aside 10% of total grant dollars for researchers who have previously not applied or had their NAIRR grant denied. Using data from the first year, the NAIRR could prioritize outreach efforts to PIs and academic institutions that did not submit a grant proposal but would be eligible for these reserved funds, and may wish to prioritize PIs from traditionally underrepresented demographics, in line with broader federal initiatives around fairness and equity.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources

In initial implementations of the NAIRR, we recommend the direction of research be led by a “bottom up” approach. By this we mean that PIs would submit proposals in areas of AI R&D that are aligned with their goals, capabilities, and expertise, rather than a government directed focus on a handful of research areas. This may encourage AI proposals from a more diverse range of study (e.g. linguistics, biology, security) and potential to develop currently unforeseeable scientific breakthroughs.

In addition to the proposal review committees, a central oversight body should be established to handle day-to-day management of the NAIRR and monitor the overall direction of AI research in the U.S. Because the NAIRR will facilitate many new AI research projects, it will be valuable to track the projects conducted in the NAIRR and to closely analyze their outcomes, especially projects that lead to subsequent beneficial social goods or which generate meaningful economic activity. This will help the government develop information to guide future funding and research efforts⁶.

2) Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Consider a Phased, Iterative Approach to the Development of the NAIRR

The Task Force may wish to consider a phased and iterative approach to the build-out and implementation of the NAIRR. In order to keep pace with the rapid acceleration of AI research and development, the Task Force should prioritize the most critical and readily-available components of a national research cloud, and make them highly usable, for the first iteration. Once the first phase of the NAIRR has been rolled out, subsequent phases can focus on the development of complementary services and capabilities for the resource.

We recommend the Task Force leverage existing infrastructure options to get off the ground quickly, meanwhile providing the research community with highly advanced technology. To train large-scale AI models, researchers must have access to a sufficient number of GPUs, CPU clusters to generate large datasets, moderately large amounts of storage to store datasets, good networking bandwidth, software to schedule jobs on the infrastructure (e.g. Kubernetes), and tooling to diagnose both hardware and software bugs. Instead of expending valuable resources developing a new ‘government cloud’ to provide these services, the NAIRR should utilize a combination of public cloud service providers and supercomputers managed by the Department of Energy National Labs (further described in Section 4, below).

High Demand for Compute

To determine technical specifications for the NAIRR, the Task Force may wish to benchmark against those from the national computing platforms of other countries. Australia’s high-performance computing (HPC) system in its National Computational Infrastructure contains a total of 155,000 CPU cores and 640 GPUs⁷, while Canada’s advanced research computing (ARC) platform provided approximately 233,000 CPU cores and 2,610 GPUs in 2021⁸. To give a sense of demand for these services, this past year Canada’s ARC platform saw its highest number

⁶ Whittlestone, J., & Clark, J. (2021). Why and How Governments Should Monitor AI Development. *arXiv*. arxiv.org/pdf/2108.12427.pdf

⁷ NCI Australia. *HPC Systems*. <https://nci.org.au/our-systems/hpc-systems>

⁸ Compute Canada. (2021). *2021 Resource Allocations Competition Results*. <https://www.computecanada.ca/research-portal/accessing-resources/resource-allocation-competitions/rac-2021-result/s/>

of applications for technical resources in its nine year history, representing a 10% year-over-year increase in proposals from 2020. The ARC, however, was only able to award 40% of the total compute requested and 22% of the total GPUs requested, indicating a growing need from researchers that outstrips current supply⁹.

Closing the compute gap between academia and industry will require ambitious technical (and financial) investment. Funding aside, a truly competitive NAIRR in the U.S. might provide 100,000 A100-equivalent GPUs and 1 million CPUs. These figures, as significant as they are, would represent an ideal research resource and one likely built well into the future. To help put these figures (and associated significant cost) in perspective, it may be useful to benchmark against leading U.S.-based digital infrastructure providers which collectively spent nearly \$97 billion on capital expenditures in 2020¹⁰.

The initial version of the NAIRR should be large enough to launch with an ability to allocate non-trivial computational and data processing resources to a handful of projects. Specifically, we imagine in its first year, the NAIRR should allocate something on the order of 30% of its computational capacity to a small number (10 or fewer) of computationally-intensive projects. Such an approach would allow researchers to ‘stress test’ the capacity of the NAIRR to facilitate large-scale experimentation, and to identify bugs in initial implementation, which could be fixed before adding further capabilities to the NAIRR. It would also differentiate the NAIRR projects from other forms of scientific research enabled by existing federal funding programs — the NAIRR would have a concrete incentive to do a set number of computationally significant experiments per year.

Prioritize Usability of Infrastructure and Research Interfaces

Access to sufficient infrastructure alone will not guarantee the success of the NAIRR; usability will also play a key role. For AI research, cloud computing systems and HPC systems aren’t always easy to use ‘out of the box’. Researchers may have to invest in building tooling to make them usable for specific scientific experiments. At Anthropic, we’ve generally found widely-available commercial clouds are easier to build software tooling on top of than HPC infrastructure. However, our investments have been non-trivial — we have multiple full-time engineers whose main job is making our infrastructure cluster stable and usable for our researchers. A successful NAIRR would need to do the same.

As a result of these tooling investments, researchers at Anthropic are able to run experiments that are both relatively easy to launch (as in it doesn’t take an individual researcher much effort to

⁹ Compute Canada. (2021). *2021 Resource Allocations Competition Results*.

<https://www.computecanada.ca/research-portal/accessing-resources/resource-allocation-competitions/rac-2021-result/s/>

¹⁰ Note that this CapEx figure includes spend on land, corporate offices, etc., in addition to cloud and data center infrastructure. Fitzgerald, C. (2021, February 5). *Follow the CAPEX: Cloud Table Stakes 2020 Retrospective*. Platformonomics.

https://www.platformonomics.com/2021/02/follow-the-capex-cloud-table-stakes-2020-retrospective/#*

boot up a large-scale multi-machine job), as well as efficient (as in the tooling helps increase the efficiency with which we utilize our cloud computing resources). These improvements ultimately help increase our experimental throughput. Researchers participating in the NAIRR may similarly need to build custom tooling to improve infrastructure usability and experiment efficiency. Those that do so should be encouraged to open source those solutions, helping other participants spend more time on research and less time on engineering needs.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

The NAIRR provides a unique opportunity to promote responsible research practices, and to gather valuable data on AI safety characteristics from aggregated research findings. These findings may be used to develop best practices in AI R&D, as well as to inform future research areas.

The NAIRR can reinforce responsible AI development by encouraging participating researchers to assess and document the broader societal impacts of their work. The Task Force may look to existing models such as the Conference on Neural Information Processing Systems (NeurIPS), an annual machine learning and computational neuroscience conference, which recently asked researchers to include a ‘broader impact’ section in submitted research papers¹¹. While the NeurIPS statements ask researchers to reflect on broader impacts at the conclusion of their research, Stanford University has developed an Ethics and Society Review board that requires researchers to consider potential negative impacts of their work as a prerequisite to receiving funding¹². Implementing similar requirements for NAIRR-enabled research may encourage participants to thoughtfully consider not only the benefits of their work, but potential harms and societal consequences, as well.

In collaboration with the research community, the government might also use the NAIRR as an opportunity to develop and iterate on assessment mechanisms to monitor the robustness, fairness, or bias of AI systems developed via NAIRR infrastructure. While evaluations to assess bias in some AI systems already exist, collaboration in a controlled research environment is likely to reveal areas of policy interest that do not have well-established benchmarks, and to yield assessment techniques that the whole community can benefit from. The government could use findings from this research conducted in the NAIRR to prioritize investment in specific areas of measurement, which has been shown to unlock follow-on work in academia and other parts of the economy¹³.

¹¹ Neural Information Processing Systems Conference. (2020, February 19). *Getting Started with NeurIPS 2020*. Medium. <https://neuripsconf.medium.com/getting-started-with-neurips-2020-e350f9b39c28>

¹² Bernstein, M. S., Levi, M., Magnus, D., Rajala, B., Satz, D., & Waeiss, C. (2021). ESR: Ethics and Society Review of Artificial Intelligence Research. *arXiv*. <https://arxiv.org/abs/2106.11521>

¹³ Whittlestone, J., & Clark, J. (2021). Why and How Governments Should Monitor AI Development. *arXiv*. arxiv.org/pdf/2108.12427.pdf

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Industry Infrastructure: Cloud Providers

Rather than a government-led effort to build net new infrastructure solely for the purposes of the NAIRR, the entity ultimately responsible for its development should leverage the advanced capabilities of industry available today — specifically, the infrastructure that has been developed relating to rentable compute and data processing systems. Utilizing the infrastructure of industry cloud providers will accelerate implementation of the NAIRR, and the NAIRR can seek to use multiple infrastructures in parallel, which will encourage competition among industry vendors during the infrastructure bidding process.

The NAIRR management entity should establish an open procurement process to solicit bids from a range of providers supplying infrastructure services and establish a panel of technical experts to review submitted proposals. The procurement terms should be clear on required functionality and desired outcomes, but otherwise remain open to a range of providers. Widening the field of available infrastructure providers also enables greater choice for researchers, allowing them to pick the provider (or providers) that most effectively meets their needs. After implementing selected providers in the first iteration of the NAIRR, the Task Force may consider a recurring procedure to solicit new bids on a regular basis (e.g. every 3-5 years) informed by the trends and feedback of researchers participating in the NAIRR.

Public Sector Infrastructure: High Performance Computing (HPC) Assets

Along with using a system based on commercial cloud technology, the NAIRR might also utilize existing and future public compute assets — specifically, supercomputers operated by the Department of Energy. The Task Force should explore repurposing these HPC assets for training AI workloads via the NAIRR. This will generate valuable information about the cost tradeoff between industry cloud infrastructure and government-owned HPC infrastructure, while also creating evidence about the costs required to make such infrastructure usable by researchers (a key success criteria of the NAIRR).

It is unknown to us whether existing HPC systems exhibit the same usability as commercial cloud computing resources; we suspect that, for frontier AI research, they don't. By evaluating how easy researchers find using HPC platforms compared to cloud platforms, the NAIRR can generate information about where to make future investments. Concretely, if it became obvious that researchers were having difficulty launching large-scale experiments on HPC infrastructure, and these experiments took longer to launch than those on commercial clouds, then that could lead to the NAIRR prioritizing investments in increasing the usability of HPC assets.

Leverage Existing Research Review Processes, While Building New Ones

As with many of the elements of the NAIRR, the research selection process represents an opportunity to build on existing programs while iterating on new parallel structures. In determining how to review and allocate funding for research proposals, the Task Force may consider a dual-track approach where a significant portion of funds (e.g. 70%) are open to PIs from any U.S. academic institution, while the remaining funds (e.g. 30%) are reserved for a small number of computationally-intensive projects. The 10% allocation for new PIs could be used for either general or computationally-intensive research proposals.

The general project track might leverage existing processes such as the NSF grant review process or DARPA's Broad Agency Announcements, with review panels composed of AI experts representative of the academic community at-large. This type of thorough and transparent review process, combined with a reserved 10% allocation for new PIs, may help distribute funding to institutions in areas that don't typically receive large-scale grants. Brookings suggests that nearly 87 metropolitan areas could be potential centers for AI R&D, on top of the 21 metropolitan areas with prominent academic institutions already receiving significant federal funding for AI projects¹⁴.

As mentioned, the NAIRR could also encourage researchers to think about bold, large-scale efforts and support them by funding a handful of computationally-intensive projects. Given the scale of these projects, PIs, academic institutions, or multiple academic institutions in collaboration would be eligible to apply. The review process for selecting these projects will likely involve an additional layer of deliberation relative to the general project grants, and could also incorporate priorities from the broader scientific community. For instance, the NAIRR could work with the NSF to identify a small number of computationally-intensive project areas, such as those that relate to AI and climate change, and then request project proposals in these areas.

In some cases, existing review processes may be too slow for promising research proposals from well-known institutions that are already leaders in AI R&D (e.g. the National Artificial Intelligence Research Institutes¹⁵). In order to accelerate their research and access to critical resources, the Task Force may consider developing a form of NAIRR access that involves minimal review for PIs who need only a modest amount of compute resources. This could further increase the impact of the NAIRR as a core piece of research infrastructure, though we believe the greatest value of the NAIRR will come from unlocking large-scale, computationally-intensive experiments.

¹⁴ Muro, M., & Liu, S. (2021, September). *The geography of AI: Which Cities Will Drive the Artificial Intelligence Revolution?* Brookings. https://www.brookings.edu/wp-content/uploads/2021/08/AI-report_Full.pdf

¹⁵ National Science Foundation. *National Artificial Intelligence (AI) Research Institutes*. <https://beta.nsf.gov/funding/opportunities/national-artificial-intelligence-research-institutes>

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Encourage Open Source Research Contributions

One failure mode of the NAIRR would be a scenario in which the U.S. government invests significant funding to democratize access to AI R&D, but the research made possible by these efforts is kept within a small cohort of actors or the discovering institution itself. To overcome this potential limitation, a significant portion of NAIRR projects should yield open source insights, either by way of software written to facilitate large-compute experimentation or published results in the form of research preprints. While we ultimately defer to the NAIRR Task Force to determine the optimal percentage of open source contributions, we recommend it be significant enough such that all participating researchers have ample work product to further assess and build upon.

Conclusion

We applaud the work of the Office of Science and Technology Policy, National Science Foundation, and the Task Force to develop a shared research ecosystem. Increasing academic access to the components of today's advanced AI technology will build on a long and successful collaboration between academia and the public sector in creating transformative technologies and advancements across the U.S. economy. We urge the Task Force to consider program design choices that leverage existing advanced infrastructure, distribute funding equitably and efficiently, and ensure the use of responsible research practices. We appreciate the opportunity to submit this response and would be delighted to answer any questions that may arise.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Argonne National Laboratory

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

RFI Response: National AI Research Resource

Input from Argonne National Laboratory

October 1, 2021

Contact: Ian Foster, [REDACTED]

Introduction

Artificial Intelligence (AI) research at Argonne National Laboratory dates back to 1963, when logician Larry Wos first started to explore the use of computers to prove mathematical theorems. Today, hundreds of Argonne scientists and engineers are engaged in the development and application of AI methods to scientific problems, as documented for example in a recent Department of Energy (DOE) AI for Science report¹. This experience has convinced us that the opportunities for application of AI methods to scientific problems are enormous and as yet largely untapped. We are thus delighted to respond to this RFI relating to a proposed National AI Research Resource (NAIRR).

The RFI requests input on many different aspects of the proposed NAIRR. We focus our response on a few areas where we believe that our experience at the intersection of AI and science allow us to make particularly pertinent and actionable recommendations

Some of the issues discussed here have previously been discussed in a Computing Community Consortium white paper².

Q1: Options to Consider

A. Goals

We propose the following goals for a NAIRR.

- Provide researchers in academia, national laboratories, and industry access to appropriately large and diverse **computing resources** to support large-scale research into both new AI methods and the innovative applications that are enabled by these new methods.
- Assemble **large, curated datasets** that, in a manner akin to ImageNet for image classification, advance the state of the art in developing AI methods and models across diverse scientific application areas.

¹AI for Science, <https://www.anl.gov/ai-for-science-report>

²A National Discovery Cloud: Preparing the US for Global Competitiveness in the New Era of 21st Century Digital Transformation, <https://arxiv.org/abs/2104.06953>

- Provide a competitive, nationwide, **peer-review process** for selecting which projects have access to these resources.
- Establish a high-quality, nationally recognized process and machinery for the systematic and sustained **intercomparison** over time of DL-based solutions to important scientific and engineering problems.

In addition to the metrics for success often used to evaluate the impacts of scientific facilities, we propose the following:

- Number of substantial AI-ready datasets assembled, and the use made of those datasets.
- Number and diversity of students and junior researchers introduced to AI methods via NAIRR.

B. Ownership and Administration

The NAIRR requires an operator with deep experience and expertise, and a deep bench of talented staff, in establishing and operating large facilities with large and diverse user communities. These factors suggest that the **Department of Energy could be the appropriate agency responsible for the implementation, deployment and administration of the NAIRR**. The DOE already has extensive, multi-decade experience in running large-scale high-performance computing centers for the national research community. More broadly, DOE has a substantial and proven experience with a sophisticated project management structure (governed by DOE Order 413.3b) that spans the lifetime of large scientific projects, including experimental facilities, large-scale computing facilities, and more recently the Exascale Computing Project. This process has been well tested and tuned over time and could be applied to run the NAIRR successfully, with appropriate controls to ensure on-time delivery within approved budget and then transition to continued successful operations. We also note that colocation with existing facilities (e.g., supercomputers) will enable new scientific methodologies such as AI-driven simulations.

D. Infrastructure capabilities

While different AI applications necessarily have varying infrastructure requirements, we wish to emphasize the needs of applications that combine AI methods for computational simulation—as when, for example, AI methods are used to guide the choice of the next simulation or to accelerate simulation kernels, or when computational simulation is used to generate training data for AI models. The growing importance of these methods in many areas of science (e.g., climate change, drug design, materials design) suggests a need for infrastructure that can **support close coupling between AI and HPC**.

E. Government Data Sets

The US government maintains many datasets of high scientific value that are, however, not easily usable by AI researchers. The reasons for this situation are varied and no single solution can overcome all obstacles. However, the NAIRR can make a substantial contribution to science by establishing methods to allow **co-location of government datasets with suitably powerful AI computers**. To give just one example, environmental researchers can, today, easily retrieve individual files from satellite imagery datasets maintained by NASA, but cannot readily train neural network models on the entirety of these large (sometimes petabyte-scale) datasets. Deployment of these datasets on storage systems colocated with AI supercomputers would allow large-scale application of AI methods to addressing important environmental questions. For other important datasets, the establishment of secure data enclaves co-located with AI supercomputers will be important for progress.

Q2: Capabilities and Services to be Prioritized

We propose two specific capabilities that we believe will have a disproportionately large impact on research in AI methods and applications.

Implement Methods for DL Model Comparison, Evaluation, and Improvement

As is well known, what gets measured gets improved. As demonstrated, for example, via the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) and the Intergovernmental Panel on Climate Change (IPCCs) Climate Model Intercomparison Project (CMIP), the establishment of codified processes, supported by appropriate tools, for comparing different approaches to a scientific or engineering question can drive great technical and scientific progress. The NAIRR can achieve broad impact on both AI research and the sciences by catalyzing and supporting similar such initiatives in areas of national priority.

To this end, we argue that NAIRR should incorporate a program that would encompass four overarching goals:

1. Development of a **general computational framework for DL model comparison, evaluation, and improvement**. This framework can include, for example, mechanisms for submitting, training, and (re)running many DL models on challenge problems, and for storing and analyzing results.
2. The **application of this framework** across specific priority research areas in close partnership with the domain science communities (e.g., cancer biology, materials discovery, synthetic biology, climate change impacts research, surrogate models).
3. **Data generation or acquisition** aimed at improving DL models in priority areas, a topic that will also couple well with research efforts in autonomous discovery. (See also below)
4. The development and incorporation into this framework of methods to improve the **interpretability, explainability, and fairness** of DL models—so that, for example, when evaluating DL models of climate change impacts, scoring metrics are evaluated with respect to their impacts on disadvantaged populations.

By promoting rigorous, data-driven comparison of alternative approaches, and enabling the systematic study of DL model effectiveness in real-world settings, this initiative will contribute to broad advances in AI methods and applications.

Create Large, AI-ready Datasets Relevant to National Priorities

The development of certain large, curated datasets such as ImageNet have been an important factor in recent AI advances. Similarly large datasets are needed to advance AI in areas of national priority, from climate science to healthcare and materials discovery. Significant effort will be needed to make such datasets AI-ready, especially with providing relevant metadata pertaining to issues such as data provenance, quality, and completeness. In many of the priority areas, sufficiently large datasets either do not exist (e.g., in materials science, despite good efforts such as the Materials Project and Materials Data Facility) and/or are not accessible in forms that permit application of AI methods (e.g., as noted above, climate science). New, substantial efforts are required to create new datasets suitable for AI research, and to deploy those datasets in ways that permit large-scale AI model training and inference.

The creation of such datasets will both drive advances in AI, by enabling AI researchers to experiment with existing and new methods in new contexts, and benefit the disciplines in which the datasets are created.

This activity will be expensive, and thus targets must be carefully chosen, with extensive community consultation. The process by which datasets are created or acquired will also require careful design and management to avoid (or at least characterize) bias.

Support the Creation of AI-for-Science Software

Modern AI benefits greatly from enormous investments by industry in powerful AI libraries and frameworks. However, these systems are inevitably designed to meet industry needs, and those needs do not always align with the needs of scientific communities. In areas where suitable software is lacking, support for its creation and maintenance will be important.

An important lesson from other scientific domains is that while graduate students can do exceptional work in exploring new methods, software produced by graduate students is rarely suitable for broad use. Building high-quality, broadly usable scientific software is hard, and requires people with specialized software engineering skills. Career paths for such people rarely exist in academic settings, but the national laboratories have a long history of constructing high-quality scientific software (e.g., mathematics libraries). Thus the construction of AI-for-science software could well be a task to be assigned to the national laboratories.

Q3: Reinforcing Ethical and Responsible AI R&D

The need for action relating to ethical and responsible AI R&D is clear. Others will surely speak to the need for robust education initiatives and for efforts aimed at creating AI-ready datasets, such as those envisioned by the recent NIH Bridge2AI program. We wish to emphasize the importance of robust and transparent processes for capturing how data are produced and curated, and models generated and shared. A useful step in this direction could be to extend the Data Management Plan that researchers must today provide with research proposals to encompass descriptions of how researchers propose to make data available with machine-readable labels or metadata.

Q4: Building Blocks

DOE national laboratories have for decades run large-scale **high-performance computing facilities** for use by the national research community, e.g., Argonne Leadership Computing Facility (ALCF), Oak Ridge Leadership Computing Facility (OLCF), and National Energy Research Scientific Computing Center (NERSC). These facilities run leading-edge supercomputers for scientific simulations, and projects are selected from an open, competitive, nationwide peer-review process. Many of these scientific applications are already starting to use AI methods, and the next-generation systems being deployed at present have been designed to support such applications and workloads, with for example 10,000s of GPUs. In other words, examples already exist of large-scale computing facilities at DOE laboratories being used for AI in science (although not exclusively). These facilities can play an important role in the NAIRR.

The DOE national laboratories also have a long history of deploying and providing community access to innovative computer systems: see for example, Argonne's Advanced Computing Research Facility³ in the late 1980s and the ALCF AI Testbed⁴ today. The national laboratories would be well

³<https://digital.library.unt.edu/ark:/67531/metadc283049/>

⁴<https://ai.alcf.anl.gov>

positioned to operate a **National AI Accelerator Laboratory** in support of NAIRR goals.

DOE national laboratories are also stewards of **important scientific datasets** that, when linked with AI computing infrastructure, can provide AI researchers with new challenges and engage them in advancing scientific goals. (To give just two examples, atmospheric radiation monitoring data in the ARM data center⁵, climate simulation datasets in the Earth System Grid Federation⁶.)

Q5: Role of Public-Private Partnerships

DOE national laboratories provide **exemplars of successful public-private partnerships in deploying leading-edge computing facilities**. Large supercomputing systems deployed at any DOE computing facility involve a multi-year partnership between the laboratory and the system vendor (in some cases, multiple vendors). Vendors are explicitly funded as part of the procurement to perform non-recurring engineering (NRE) activities to meet the ambitious goals of the system. These NRE components are often essential to ensure that the technologies needed to deliver the system are available in the timeframe needed for the system deployment. In other words, NRE accelerates the availability of new technologies in a vendor's product. Furthermore, DOE has a history of separately funding vendors to develop new technologies independent of any specific procurement. Examples of such programs include FastForward, DesignForward, and PathForward, each which funded multiple vendors to accelerate their technology roadmaps to enable exascale computing. As a result, three exascale systems will soon be available at Oak Ridge, Argonne, and Lawrence Livermore National Laboratories.

Q6: Possible Limitations to Democratization

Design resource allocation policies to promote democratization

Methods for allocating scarce scientific resources that are based purely on merit-based peer review can be exclusionary due to winner-takes-all effects. Experience with national scientific facilities such as supercomputers emphasize the importance of also implementing policies designed to combat that effort. Specifically:

1. Set aside a substantial chunk of the resources to explicitly address fairness issues in access. Thus, for example, any researcher or student from an educational institution might be able to obtain some limited "startup" allocation.
2. Also allocate substantial resources to provide help to those who need it. Online help, training courses, etc.
3. And support a substantial outreach program to engage institutions and communities who are not currently active in the use of the resource.

The latter two activities, in particular, are not cheap (it takes much more human labor to support 10,000 rather than 100 users), but are essential if the NAIRR is to engage a broad community.

⁵<https://www.arm.gov/data/>

⁶<https://esgf.llnl.gov>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

ACM US Technology Policy Committee

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1, 2021

**COMMENTS IN RESPONSE TO RFI ON AN IMPLEMENTATION PLAN
FOR A NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH RESOURCE
(DOCUMENT NUMBER 2021-15660)**

The non-profit [Association for Computing Machinery](#) (ACM), with more than 50,000 U.S. members and approximately 100,000 worldwide, is the world's largest educational and scientific computing society. ACM's [US Technology Policy Committee](#) (USTPC), currently comprising more than 140 members, serves as the focal point for ACM's interaction with all branches of the US government, the computing community, and the public on policy matters related to information technology. As such, the Committee strives to serve as an apolitical source of expert information.¹

Overview

USTPC is pleased to respond to the Request for Information in the above-referenced proceeding² issued jointly by the National Science Foundation and White House Office of Science & Technology Policy.³ Before doing so in detail below, however, the Committee wishes to underscore its fundamental recommendation that the new National Artificial Intelligence Research Resource (NAIRR) should be undertaken *only* if it is independently funded. Creation and operation of the NAIRR thus should not reduce current levels of federal financial support for Artificial Intelligence (AI) research by any other arm of the government.

¹ To arrange for a technical briefing from USTPC and other ACM expert members, please contact Adam Eisgrau, ACM Director of Global Policy & Public Affairs, at acmpo@acm.org or 202-580-6555.

² See 86 FR 39081 (July 23, 2021).

³ Principal authors of these Comments for USTPC were: Jim Kurose, Distinguished Professor in the College of Information and Computer Sciences and Associate Chancellor for Partnerships and Innovation at the University of Massachusetts Amherst; and Arnon Rosenthal of Bedford, Massachusetts, a consultant and researcher specializing in data integration, databases, access controls, and their interactions with adjacent technologies. Also contributing were: ACM Technology Policy Council Chair Jim Hender, USTPC Chair Jeremy Epstein, Vice Chair Alec Yasinsac, and USTPC members Joshua Kroll and Bulent Yener.

Q1. What options should the Task Force consider for any of roadmap elements A - I?

Element C —A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources

A common infrastructure for all. Currently, computing research is funded at the national level by a relatively few agencies each (understandably) funding access to research computing infrastructure primarily for the researchers that it directly sponsors, reflecting the culture and organization of the funding organization.⁴ Thus, for example, while some DOE-funded researchers not supported directly by NSF are able to access NSF-funded resources, and *vice versa*, such resource sharing is not common. As a national resource targeted at research in a specific area, the NAIRR should not be partitioned or “siloeed” according to funding agency.

Coupling NAIRR allocation review with agency review. The demand for NAIRR resources is virtually certain to exceed their availability. Accordingly, a process to assess the readiness and appropriateness of proposed resource allocations will be required. Several national models exist for accomplishing such resource allocation⁵ that have generally served the community well and profitably might be emulated by the NAIRR. USTPC strongly recommends, however, that NSF’s specific practice of de-coupling approval of basic research decisions⁶ from others made in the LRAC/XRAC-resource allocation process *not* be adopted. These two sets of decisions, in USTPC’s view, ideally should be made in tandem, thus permitting the totality of a project to be fully and accurately evaluated.

Obtaining maximum value from data. Data have fundamental different characteristics, use patterns, and cost structures than the “use it or lose it” characteristics of computational resources. Unlike processors, data can be accessed and used simultaneously by many recipients with minimal marginal costs; “public” datasets can also be downloaded and processed with users’ own computational resources. If NAIRR is hosted on a cloud, processing or export might be done on the user’s cloud account, sparing NAIRR the burden of billing. It also should be possible, and would certainly be desirable, for the NAIRR to negotiate favorable rates for data export to other clouds. It would be valuable for the NAIRR to support these, and other, approaches towards data access, migration, and use.

⁴ The national computing research infrastructure is presently funded primarily by the Department of Energy and National Science Foundation with smaller scale investments being made by the National Oceanic and Atmospheric Administration, National Institute of Health, NASA, and other agencies.

⁵ These include: the Leadership Resource Allocation Committee (LRAC) and XSEDE Resource Allocation Committee (XRAC) for computational resources funded via the NSF Office of Advanced CyberInfrastructure (OAC) and INCITE for resources funded via the DOE Office of Science in the U.S. Department of Energy. LRAC and XRAC, which provide merit-based, peer-review allocation processes for resources that would otherwise be oversubscribed.

⁶ For example, traditional agency-convened merit-review panels make recommendations to fund research that requires access to computational resources but do not provide funding for access to these resources.

A community-based technical advisory committee. As experience is gained once the NAIRR becomes operational, USTPC recommends that lessons to be learned and efficiencies gained from that experience might best be analyzed by a computing community-based expert advisory committee charged broadly with providing the NAIRR with advice and feedback. By no means incidentally, USTPC also notes that the composition of such a community-based committee should be carefully and deliberately defined to ensure democratized access to the NAIRR’s resources by diverse researchers and for diverse socially relevant purposes. Care thus also should be taken to assure that the committee’s makeup broadly reflects both the diversity of the research community and the communities intended to benefit from funded research.

Element H — A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector.

Coordinated decision making. Many of the AI researchers utilizing the NAIRR are likely either to be funded by one or more federal agencies or work at a federal agency or laboratory. As noted above, it is important that NAIRR allocation decisions be made in concert with research funding allocation decisions made by these agencies.

Overhead expense relief. USTPC recommends that NAIRR access *not* be subject to overhead recovery costs if such access is explicitly allocated in an agency research award. As noted in the Report of NSF’s workshop on *Enabling Computer and Information Science and Engineering Research and Education in the Cloud*,⁷ this would “level the playing field” between an organization’s option to (i) acquire hardware and software computing research assets, which are not subject to overhead, and (ii) acquiring access to similar assets in the cloud or elsewhere, which typically is subject to overhead charges at most institutions.

Element I — Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.⁸

Multi-agency oversight and budget. Researchers funded by Federal agencies, or those working within them, will likely be among those with access to the NAIRR. Thus, USTPC suggests, multiple agencies should be assigned an oversight and governance role within the NAIRR. Several agencies with significant investment in AI research (e.g., NSF, DOE, NIH, NIST, NOAA) might appropriately be considered “lead” agencies in such a governance structure. Given AI’s pervasive and growing application in many spheres, however, it also will be important for other agencies (e.g., addressing Agriculture, Commerce, Transportation, Justice, and more) to play roles in defining and operationalizing the NAIRR.

⁷See www.researchgate.net/publication/329911307_Enabling_Computer_and_Information_Science_and_Engineering_Research_and_Education_in_the_Cloud (July 2018).

⁸ USTPC notes that much of the following material also potentially is germane to implementation roadmap element B (“A plan for ownership and administration” of the NAIRR) and requests that it also be considered in that context.

Budget flow and agencies. If the budget for the NAIRR is funneled through a single agency, that agency's researchers (whether externally funded or in-house) may be perceived as having priority access to its resources, even if this is not the case in fact. Without budget flow, a resource like the NAIRR may become less visible within the agency. To avoid such an appearance of conflict of interest, and to ensure agency engagement, NAIRR should be structured to route funding through multiple agencies without, USTPC again emphasizes, being required to draw upon their existing budgets. Potential models for such multi-agency funding include major research instrumentation projects, such as telescopes and accelerators.⁹ A non-agency, non-governmental entity may serve to more easily and effectively manage the NAIRR, and as an added benefit help facilitate the development of public-private partnerships.

Funding scope. The scope of the NAIRR's support must be carefully defined and tailored to AI *research computing infrastructure* costs rather than to other AI research components. Expenses such as faculty and graduate student compensation, travel, and more all are necessary and complementary to AI research computing expenses. NAIRR funding, however, will be most effective if dedicated to covering purely AI research computing infrastructure costs.

Open research. Similarly, it should be made clear that the NAIRR's principal aim is enabling open, basic AI research. Funding for computational resources for the commercial development of AI technologies should be provided by other means, such as venture capital investment, economic development fund underwriting of early-stage technologies, or direct industry investment in established technologies by more mature companies.

Q2. Which capabilities and services provided through the NAIRR should be prioritized?

The inquiries made by the RFI in this proceeding clearly are predicated on the authoring agencies' conviction that the NAIRR should support much more than mere access to computational capacity and raw data sets. USTPC broadly concurs with this fundamental premise, and with regard to the following specific capabilities and services:

Computing. Whenever possible NAIRR should leverage commercial cloud offerings, developing its own specialized computing resources only after conducting a comprehensive cost/benefit analysis that justifies doing so. The NSF's 2018 NSF workshop Report on *Enabling Computer and Information Science and Engineering Research and Education in the Cloud* cited at note 7 above provides a thoughtful discussion of the many advantages (and possible disadvantages) of this approach.¹⁰

⁹ Examples include AURA, which operates and builds world-class astronomical observatories for NSF and NASA, and US IGNITE, which helps manage NSF's PAWR initiative.

¹⁰ See also "[Cloud Access for NSF CISE Research](https://cra.org/cloud-access-for-nsf-cise-research/)" at <https://cra.org/cloud-access-for-nsf-cise-research/>.

Data. To maximize access to data and its utility to researchers,¹¹ USTPC suggests that the NAIRR:

- recognize that while money alone can provide access to computing capacity, it cannot always guarantee access to data necessary for a particular research effort.¹² Data thus should be presumed to be a community resource to be managed and made widely accessible accordingly;
- enhance the value of communal data sets by encouraging that user commentary and experiences be associated with all accessible data sets;¹³
- maximize the ways in which all accessible data sets are characterized, categorized, and cross-referenced to enable users to most easily interpret the data, determine its fitness for use, and assess whether and how it may be biased;
- mandate that all accessible data be accompanied by metadata sufficient to understand its provenance, meaning, and trustworthiness;
- assure that the collection of NAIRR-hosted datasets is “FAIR”: Findable, Accessible, Interoperable, and Re-usable to the maximum (and likely varying) degree to which each such goal is feasible;
- foster interoperability by devising and encouraging data generators to adopt standards for its form, metadata, and organization but, to avoid discouraging contributions, nonetheless accept data in native form if augmented by sufficient descriptive material to permit others to supplement or complete its description;¹⁴
- develop the capacity to host anonymized or otherwise “privatized” datasets (created, for example, using differential privacy) derived from confidential origin datasets, and to provide information describing how these data sets were created and might be utilized. The NAIRR also should consider hosting non-public datasets through an interface designed to provide only “sanitized” products (*e.g.*, selected “safe” statistics, not original data);

¹¹ While pursuing and sharing data is a valuable goal, *data alone does not guarantee progress in AI research*. Prioritizing “more data” too often drives policy making when “better data” or “understandable data” should be preferred.

¹² While a great deal of data (often in huge sets) is generated by and accessible within large corporations and other private entities, clearly not all researchers have access to these resources. To the extent possible (*e.g.*, as constrained by competitive commercial concerns of data owners, and privacy or security considerations) the NAIRR can and should make such data broadly available to researchers everywhere.

¹³ For example, “Q&A” documentation and shared code (*e.g.*, as currently available through www.stackoverflow.com) would enhance data set value for all researchers.

¹⁴ While this objective might theoretically best be achieved through mandated standards compliance (or by declaring parts of a dataset to be semantically the “same as” part of another set), USTPC is aware that many data owners may be unwilling or unable to conform to such a standard.

- host community-developed code designed to enhance and enable the manipulation of datasets, as well as models built on such information;¹⁵
- broadly define the term “data resources” to include “higher-level” data products. Such synthesized data structures, like knowledge graphs, are critical for AI research but often are available at scale only within industry settings.¹⁶

People. Raw computational capabilities and datasets are a necessary, but far from sufficient, set of resources to meet the NAIRR’s goals. Perhaps most importantly, people-centric services like education, training, expert consulting, and community outreach¹⁷ all must be key components of an accessible NAIRR that seeks to successfully serve a broad and diverse community of AI researchers.

Q3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Democratizing access: beyond supporting agency-funded AI research. USTPC recommends that, as NSF presently does in connection with funding computational resources,¹⁸ the NAIRR set aside a specific and not insignificant percentage of its total available resources for otherwise unfunded developmental, exploratory, and other meritorious projects that demonstrate the need and appropriateness for NAIRR resources. A portion of these reserved resources should in USTPC’s view be dedicated to projects that concern issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability.

USTPC also reiterates here its recommendation responsive to Question 1, Element C above that NAIRR be guided in part by a diverse community-based committee with a broad community perspective (in addition to XRAC/LRAC/NSERC-like bodies) charged with advising on and overseeing resource allocations. The committee should be expressly charged with, among other responsibilities, ensuring democratized access by diverse researchers to the NAIRR’s resources. As also noted above, the committee’s makeup also should broadly reflect both the diversity of the research community and the communities intended to benefit from funded research.

¹⁵ More generally, a model “commons,” similar to NIH’s notion of a “data commons,” might include systems like the Generative Pre-trained Transformer 3 (GPT-3) for use by the research community.

¹⁶ See, e.g., “Open Knowledge Network” NITRD Big Data IWG workshop report (November 2018).

¹⁷ The NAIRR might consider, in particular, outreach by distributed teams of local “champions” modeled, perhaps, XSEDE “campus champions”.

¹⁸ In the case of NSF-funded computational resources, 20% are reserved for specific uses not subject to LRAC/XRAC review, or to be awarded at the discretion of the Director of the institution operating that resource.

Finally in this context, USTPC notes that the Association for Computing Machinery's Diversity, Equity, and Inclusion Council also could be an important resource for the NAIRR in populating the recommended committee and developing practices to improve diversity in NAIRR's organization and programs by leveraging ACM's broad and deep connectivity to the nation's AI research community.

Q4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

USTPC recommends that the NAIRR consult existing recent pilot existing programs and resources relevant to the NAIRR's vision and mission, such as NSF Cloudbank, NIH STRIDES, and the 2018 NSF Workshop Report on *Enabling Computer and Information Science and Engineering Research and Education in the Cloud*.

Q5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

USTPC sees significant potential for NAIRR to maximize the impact of its resources by catalyzing a public-private partnership¹⁹ to collaboratively construct the NAIRR's physical infrastructure in the cloud with access afforded by commercial cloud vendors. Such a structure would allow the NAIRR itself to focus on what may well be its "highest" functions: providing coordination and funding for resource access and use. It also might productively explore offering any of a myriad of services that may be "layered" on top of the physical infrastructure, much as Cloudbank and STRIDES, both referenced above, do so now.²⁰

Conclusion

USTPC commends the NSF and OSTP for initiating this proceeding to help launch the NAIRR which, USTPC believes, can bring significant benefits to the AI research community and the nation. As noted in the updated *2019 National AI R&D Strategic Plan*,²¹ AI holds tremendous promise for society, and provides important opportunities for national economic competitiveness, health and welfare, defense, and security. The NAIRR can play an important role in achieving these benefits by providing democratized access to AI research computing and data infrastructure, and to services built on and around this infrastructure. USTPC looks forward to working with NSF and OSTP to those vital ends.

¹⁹ These Comments also touch upon private-public partnership potential in the context of responding to RFI Q1, Element H above.

²⁰ This concept is not theoretical. Multiple cloud providers have collaborated to provide resources to researchers funded under the NSF BIGDATA program and commercial cloud services also have hosted large NOAA datasets.

²¹ See <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019>.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Atlantic Council GeoTech Center

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Response to NSF/OSTP RFI
Submitted by the Atlantic Council GeoTech Center
September 27, 2021

The RFI asks for responses to the following six questions.

1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

Enhanced trust and confidence in the digital economy is founded upon personal privacy, data security, accountability for performance and adherence to standards, transparency of the internal decision-making algorithms, and regulations and governance for digital products and services. Trust and confidence in the digital economy is diminished by practices that do not protect privacy or secure data, and by a lack of legal and organizational governance to advance and enforce accountability.¹ Data breaches, malware embedded in downloaded apps, unfiltered mis- and disinformation, and the lack of governance models to effectively address these harms all contribute to the degradation of social and civic trust. This degradation undermines economic and civic confidence, is costly,² constrains the growth of the digital economy,³ and has destabilizing effects on society, governments, and markets. Trust and confidence in the digital economy is essential for open societies to function, and for resilience against cascading effects of local, regional, or national economic, security, or health instabilities.

In summary, a key goal of the NAIRR testbed is to ensure, as AI technologies and applications are developed, that the future digital economy—powered increasingly by AI—is trusted, is trustworthy, and that it protects individual privacy and rights. The following table provides specific options for several of the roadmap elements. Each is discussed further in the subsequent sections of this paper.

¹ Amon, “Toward a New Economy of Trust.”

² World Economic Forum, “Why trust in the digital economy is under threat,” accessed March 26, 2021, <http://reports.weforum.org/digital-transformation/building-trust-in-the-digital-economy/>, citing an estimate by McAfee that the costs associated with cybersecurity incidents approximated \$575 billion in 2014; Accenture, *Securing the Digital Economy: Reinventing the Internet for Trust*, 16, accessed March 26, 2021, https://www.accenture.com/us-en/insights/cybersecurity/_acnmedia/Thought-Leadership-Assets/PDF/Accenture-Securing-the-Digital-Economy-Reinventing-the-Internet-for-Trust.pdf#zoom=50. Cites five-year loss of foregone revenue from 2019 to 2023 to be \$5.2 trillion, calculated using a sample of 4,700 global public companies.

³ Congressional Research Service, *Digital Trade and U.S. Trade Policy*, 11, May 21, 2019, accessed March 26, 2021, <https://crsreports.congress.gov/product/pdf/R/R44565>; Alan B Davidson, “The Commerce Department’s Digital Economy Agenda,” Department of Commerce, November 9, 2015, accessed March 26, 2016, <https://2014-2017.commerce.gov/news/blog/2015/11/commerce-departments-digital-economy-agenda.html>. Davidson identifies four pillars: promoting a free and open Internet worldwide; promoting trust online; ensuring access for workers, families, and companies; and promoting innovation.

Roadmap Element	Options
A. Goals	Emphasize trust and trustworthiness of the digital economy as one develops AI technologies and applications. Make this a central tenet of the testbed.
B. Ownership and administration	Construct a national coordinating office to lead and manage the testbed. This will include public- and private-sector representatives, state and local governments, and international outreach.
C. Governance	Establish an advisory council with representation of all stakeholders. A leadership council would be the primary decision-making group, representing the primary funding sources.
D. Capabilities	Use as a model the secure, distributed, cloud-based information environments developed for the Intelligence Community.
E. Barriers to use of data	Emphasize the need to achieve user-acceptable levels of privacy and trust in how the data are used and how individual rights are preserved as AI-based applications become more widespread.
F. Security requirements	Adhere to federal standards and guidelines for trusted information technology, but with up-to-date implementation practices.
G. Privacy and civil rights	Include privacy-preserving technologies in the priority research goals of the testbed and in demonstrations with real users.
H. Sustainment	Include international partners in the design and evaluation of the testbed and the selection of research agendas. This may include both the private sector and with other nations.
I. Agency roles	This is a national-scale initiative, requiring broad involvement by the federal, state, and local governments. The many stakeholders (owners of data, lead agencies for problem areas, funders of research) all must be involved in the long-term sustainment and management of the testbed.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

The following discussion helps inform the following NAIRR topics:

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research.

To enhance trust and confidence in artificial intelligence and other digital capabilities, technologies must objectively meet the public’s needs for privacy, security, transparency, and accountability.

The growth of digital economies is changing how trust is valued by institutions, businesses, and the public.⁴ The traditional view of trust is expressed in terms of the security of a business transaction. The increase in cyberattacks, identity theft, social media disinformation campaigns, and the use of autonomous decision-making software, introduces new factors that affect trust. Trust in a firm’s reputation and ethical practices, privacy protection, and how personal data are used depend on technology, business practices, and the public’s perception of how well these components of trust are protected.

Not everyone has the same perception of what is trustworthy. However, reaping the benefits of the digital economy requires a high level of trust among users. Therefore, government and industry should work to enhance the transparency and accountability of digital systems to improve trustworthiness. Challenges include the following: (i) views on personal privacy protection are context-dependent, vary by culture or location, and may be formalized in different terms across nations, regions, and states; and (ii) as automated decision-making algorithms proliferate, new applications reveal trust weaknesses regarding implicit bias, unethical use of personal data, and lack of identity protection.

Trustworthiness needs to be prioritized and empirically demonstrated in the evolving market. Building trust involves educating all participants on the fundamental value of trust in the digital economy and ensuring digital systems reflect individual and societal conceptions of trust. There must be national and international standards for judging how well technologies and systems protect trust. Professional organizations that audit for trust in the digital economy will strengthen accountability.

*Finding: The European Union’s General Data Protection Regulation uses data protection rules as a trust-enabler.*⁵

As European Union (EU) member nations work to conform national rules and laws to the General Data Protection Regulation (GDPR), the European Commission notes that these steps may strengthen trust relationships. Other nations propose that a global framework for cross-border Internet policies may be able to protect data security and privacy while still allowing national laws and regulations as a part of the approach if certain trust relationships are maintained. For both approaches, a set of rules or principles provides the foundation for trust.

⁴ Frank Dickson, “The Five Elements of the Future of Trust,” IDC, April 22, 2020, accessed March 26, 2021, <https://blogs.idc.com/2020/04/22/the-five-elements-of-the-future-of-trust/>.

⁵ “Communication from the Commission to the European Parliament and the Council. Data protection rules as a trust-enabler in the EU and beyond – taking stock,” COM/2019/374 final, European Union, July 24, 2019, accessed March 26, 2021, <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=COM:2019:374:FIN>.

The GDPR⁶ establishes regulations for data security and privacy that apply to any organization that collects or uses data related to people in the EU. The entire data chain is covered by the GDPR, including data collection, processing, storing, and managing.

The GDPR comprises principles that govern data protection and accountability for those who process data. There are technical measures for data security, and organizational design principles for data protection. Data privacy is expressed in terms of privacy rights, including the right: to be informed, to rectification, to erasure, to restrict processing, to data portability, and to object, and the right of access. There are also rights in relation to automated decision-making and profiling. The governance mechanism centers on Data Protection Authorities that work to align each EU member nation's approach to data security and privacy to conform with the GDPR. These Data Protection Authorities have enforcement powers and the ability to levy fines when a GDPR rule is violated.

*Finding: Current approaches to machine learning and big data analytics risk weakening data protection rules.*⁷

Data privacy protection is vulnerable to advanced data analytics that can infer personal identifiable information by joining loosely related data sources. As a result, the growing use of current machine learning methods applied to large, multi-source data sets highlights potential limitations in the GDPR where such computational methods can infer data originally made private. The development of new data science capabilities may require research on new privacy-preserving technologies for nations to remain compliant with the GDPR. With increasing amounts of personal medical and genetic information being held in data repositories, this need is urgent.

Finding: Evolving US data privacy approaches consider outcome-based methods, versus prescriptive methods.

The development of data privacy laws in the United States is an evolving patchwork, with more than one hundred and fifty state data privacy laws proposed in 2019.⁸ There is no overall federal data privacy law.

⁶ "General Data Protection Regulation," Intersoft Consulting, <https://gdpr-info.eu/>.

⁷ T. Timan and Z.Á. Mann, eds., *Data protection in the era of artificial intelligence. Trends, existing solutions and recommendations for privacy-preserving technologies*, Big Data Value Association, October 2019, accessed March 26, 2021, https://www.bdva.eu/sites/default/files/Data%20protection%20in%20the%20era%20of%20big%20data%20for%20artificial%20intelligence_BDVA_FINAL.pdf.

⁸ "2019 Consumer Data Privacy Legislation," National Conference of State Legislatures, January 3, 2020, accessed March 26, 2021, <https://www.ncsl.org/research/telecommunications-and-information-technology/consumer-data-privacy.aspx>.

One instance of federal legislation for data privacy proposed in the 117th Congress⁹ includes the following key privacy features, which are viewed as outcome-based.¹⁰

- Transparent communication of the privacy and data use policy
- Affirmative opt-in and opt-out consent
- Preemption, in which the proposed statute would preempt most state laws with limited exceptions for data breaches, and other limited situations
- A right to action, enforced at the federal or state level, to address alleged violations
- Independent audit of the effectiveness and appropriateness of the privacy policy for each entity providing data services

The National Institute of Standards and Technology (NIST) Privacy Framework describes a risk- and outcomes-based approach to establishing privacy protection practices in an organization. Organizations can vary the technologies and design of the privacy protection aimed at satisfying performance outcomes. This may be advantageous when the technologies and applications are changing at a fast pace, e.g., artificial intelligence (AI) and the Internet of Things (IoT).¹¹

Finding: New information technologies compel automated compliance testing.

New information technologies and advanced data capabilities challenge current methods of compliance and enforcement. The variety of new ways to collect, process, and analyze data is increasing at a fast rate, while compliance often is determined on a case-by-case basis by regulatory and legal experts. To keep pace, automated testing for compliance with data privacy regulations is necessary.

Table 1 portrays some of the challenges and solutions for achieving automated compliance testing. This research agenda identifies the following key developments: standards, new privacy-preserving technologies, and automated methods to establish compliance. Privacy-preserving technologies are an active research area, and include the following: secure multiparty computation,

⁹ “Information Transparency and Personal Data Control Act,” fact sheet, accessed March 26, 2021, https://delbene.house.gov/uploadedfiles/delbene_consumer_data_privacy_bill_fact_sheet.pdf; Information Transparency & Personal Data Control Act, H.R. 2013 — 116th Congress (2019-2020), accessed April 2, 2021, https://delbene.house.gov/uploadedfiles/delbene_privacy_bill_final.pdf.

¹⁰ “Developing the Administration’s Approach to Consumer Privacy,” *Federal Register*, September 26, 2018, accessed March 26, 2021, <https://www.federalregister.gov/documents/2018/09/26/2018-20941/developing-the-administrations-approach-to-consumer-privacy>; Alan Charles Raul and Christopher Fonzone, “The Trump Administration’s Approach to Data Privacy, and Next Steps,” Sidley Austin LLP, October 2, 2018, accessed March 26, 2021, <https://datamatters.sidley.com/the-trump-administrations-approach-to-data-privacy-and-next-steps>.

¹¹ National Institute of Standards and Technology, “NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management, Version 1.0,” January 16 2020, accessed March 26, 2021, https://www.nist.gov/system/files/documents/2020/01/16/NIST%20Privacy%20Framework_V1.0.pdf.

(fully) homomorphic encryption, trusted execution environments, differential privacy, and zero-knowledge proofs.

Table 1. Big Data Value Association Strategic Research and Innovation Agenda

Challenges	Solutions
A general, easy-to-use, and enforceable data protection approach	Guidelines, standards, law, and codes of conduct
Maintaining robust data privacy with utility guarantees	Multiparty computation, federated learning approaches, and distributed ledger technologies
Risk-based approaches calibrating data controllers' obligations	Automated compliance, risk assessment tools
Combining different techniques for end-to-end data protection	Integration of approaches, toolboxes, overviews, and repositories of privacy-preserving technologies

Source: Timan and Mann 2019¹²

The value of privacy-preserving technologies involves trade-offs between privacy and utility—how useful is the resulting data—both of which are context dependent.¹³ Affecting these trade-offs are the technical methods, the technical definitions of privacy, and the specifications of the privacy laws. The technical methods (e.g., anonymization, sanitization, and encryption) operate on data in different ways. The technical definition of privacy varies by application and the user's perceptions of risk versus the benefit of making personal data available. Privacy laws vary across nations, challenging the uniform application of technical methods. For both professionals and members of the public, making trade-offs between privacy and utility remains challenging. This is partially due to the absence of definitions of and standards for measuring privacy and the social benefits obtained from making data available for use by others.

Priority: Trust and confidence in digital capabilities requires businesses and governments to focus on the responsible use of technology.

Increasing trust and confidence in emerging technologies, such as AI, requires a recognition by both businesses and governments that they have an obligation to use technology responsibly, ensuring that technology has a positive impact on society, especially with regards to equality and

¹² Timan and Mann, *Data protection*.

¹³ Daniel Bachlechner, Karolina La Fors, and Alan M. Sears, "The Role of Privacy-Preserving Technologies in the Age of Big Data," proceedings of the 13th Pre-ICIS Workshop on Information Security and Privacy, San Francisco, December 13, 2018, accessed March 26, 2021, https://www.albany.edu/wisp/papers/WISP2018_paper_11.pdf; Felix T. Wu, "Defining Privacy and Utility in Data Sets," *University of Colorado Law Review* 84 (2013), accessed March 26, 2021, http://lawreview.colorado.edu/wp-content/uploads/2013/11/13.-Wu_710_s.pdf.

inclusion.¹⁴ Developing and innovating responsibly means ensuring that (i) ethical frameworks and policies exist to guide organizations during all aspects of a product’s development and deployment, (ii) fairness in design is emphasized from the outset, and that (iii) questions around the manner in which technologies will be used are given the same rigorous examination as technical issues. As technological capabilities evolve and become more deeply intertwined in all aspects of society, businesses and governments must put ethics at the center of everything they do.

Priority: Build in trust-enabling technologies, measure performance against standards, conduct independent compliance audits.

The digital economy relies on achieving a high level of trust and confidence on a continuing basis as technologies evolve. Trust and confidence-enabling technologies must be developed and built into the components of the digital economy infrastructure; a detailed understanding of the trade-offs between privacy versus utility is an essential foundation. Such technologies must be paired with similar civic norms, practices, and rules designed to enhance confidence in the digital economy. To assure businesses that they remain compliant with data protection regulations as they modernize their practices, automated compliance testing, accompanied by standards of performance, is needed. To establish transparency for automated decision-making algorithms, standards for the measurable performance, i.e., the output results, are necessary. Independent assessments of the compliance testing and algorithmic transparency by professional auditing organizations could enhance trust among all participants in the digital economy and aid accountability and governance; such methods should be explored. However, mechanisms for compliance testing and auditing by regulators are also necessary.¹⁵

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Recommendation: Assess standards relating to the trustworthiness of digital infrastructure.

Congress should direct an assessment by the National Academies of Sciences, Engineering, and Medicine of the current national and international standards relating to the trustworthiness of digital infrastructure to support the digital economy. “Trustworthiness of an information system is defined as the degree to which an information system (including the information technology components that are used to build the system) can be expected to preserve the confidentiality,

¹⁴ Kirsten Martin, Katie Shilton, and Jeffrey Smith, “Business and the Ethical Implications of Technology: Introduction to the Symposium,” *Journal of Business Ethics* 160, 307–317 (2019), accessed April 16, 2021, <https://doi.org/10.1007/s10551-019-04213-9>

¹⁵ Nicholas Confessore, “Audit Approved of Facebook Policies, Even After Cambridge Analytica Leak,” *New York Times*, April 19, 2018, accessed March 26, 2021, <https://www.nytimes.com/2018/04/19/technology/facebook-audit-cambridge-analytica.html>.

integrity, and availability of the information being processed, stored, or transmitted by the system across the full range of threats.”¹⁶

Due to the increasing complexity of the digital infrastructure, the assessment should also review design standards for complex systems-of-systems from the perspective of trustworthiness. The overall assessment focuses on systems that support the digital economy. The study should assess the sufficiency of existing standards to guide improvements in trustworthiness, identify where new standards are needed, and recommend the data collection and testing methods that would enable ongoing assessments.

Recommendation: Produce a framework for assessing ethical, social, trust, and governance considerations associated with specific current and future use cases for AI.

The administration should request the National Academy of Sciences to produce a framework for assessing ethical, social, trust, and governance considerations associated with specific current and future use cases for AI solutions. The framework should identify where new federal standards and rules are needed. This guidance should be developed with the participation of relevant executive branch departments and agencies, and in consultation with private industry, academia, members of the public, and government and industry representatives from foreign partners.

Recommendation: Educate the public on trustworthy digital information.

Congress should establish a grant program led by NSF for the purpose of developing a curriculum on trustworthiness of information—distinct from the trustworthiness of information systems—in the digital age. This curriculum should be created by a consortium headed by a university or coalition of universities. The program should be administered by select universities, with the participation of US information providers. The goal should be to educate the public on how to assess the trustworthiness of information—its credibility, truthfulness, and authenticity, and to develop tools that students and members of the public can use and benefit from on a regular basis.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Recommendation: Create measurement methods and standards for evaluating trust in the digital economy.

The administration should direct the National Institute of Standards and Technology (NIST) to establish methods for evaluating users’ trust in the digital economy given the increasing use of AI, big data analytics, and automated decision-making algorithms. This work builds on the Commission on Enhancing National Cybersecurity’s *Report on Securing and Growing the Digital*

¹⁶ National Institute of Standards and Technology, *Security and Privacy Controls for Information Systems and Organizations*, Special Publication 800-53, Revision 5, September 2020, accessed April 16, 2021, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf>.

*Economy*¹⁷ and the *National Strategy for Trusted Identities in Cyberspace*.¹⁸ One assessment framework example¹⁹ describes measures of: “(i) user trust in the digital environment, e.g., data privacy, security, private sector efforts to control the spread of misinformation, and private sector adherence to cybersecurity best practices; (ii) the user experience, i.e., the effort needed to interact with the digital environment; (iii) user attitudes, e.g., how trusted are government and business leaders; and (iv) user behavior, i.e., how much do users interact with the digital environment.”

The administration should create a coalition to develop international standards for achieving trust in the digital economy. The coalition should include representatives from NIST, the Federal Trade Commission (FTC), private industry, Federally Funded Research and Development Centers (FFRDCs), University Affiliated Research Centers (UARCs), and international standards organizations. The United States and like-minded nations and partners should develop national assessments of trust in the digital economy using these standards.

Recommendation: Empower an organization to audit trust in the digital economy.

Congress should establish or empower an organization to audit the efficacy of measures designed to ensure trust in the digital economy and assess conformance to current and future standards designed to enhance and maintain such trust. Independent third parties or the Government Accountability Office (GAO) are examples of where such auditing organizations could be housed.

As part of this process, the auditing organization could provide recommendations to Congress on legislation that would enhance existing trust measures, develop new trust measures, and create trust performance standards. The auditing organization should also provide a mechanism through which the public and industry can raise topics and concerns for attention and, for cases where assessments or audits were done, include an ombudsman function for assessment appeals, identification of new information, or adjudication of concerns in a manner distinct from political influence.

The administration should work to establish a similar auditing program with EU members of the International Organization of Supreme Audit Institutions.

¹⁷ Commission on Enhancing National Cybersecurity, *Report on Securing and Growing the Digital Economy*, December 1, 2016, accessed March 26, 2021, <https://www.nist.gov/system/files/documents/2016/12/02/cybersecurity-commission-report-final-post.pdf>.

¹⁸ White House, “National Strategy for Trusted Identities in Cyberspace, Enhancing Online Choice, Efficiency, Security, and Privacy,” April 2011, accessed March 26, 2021, https://obamawhitehouse.archives.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf.

¹⁹ Bhaskar Chakravorti, Ajay Bhalla, and Ravi Shankar Chaturvedi, “How Digital Trust Varies Around the World,” *Harvard Business Review*, February 25, 2021, accessed April 16, 2016, <https://hbr.org/2021/02/how-digital-trust-varies-around-the-world#:~:text=To%20that%20end%2C%20in%20partnership,user%20experience%3B%20the%20extent%20to.>

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Recommendation: Conduct demonstration projects involving artificial intelligence to improve delivery of public- and private-sector services at local, state, and federal levels.

Congress should authorize and appropriate funds for AI demonstration projects that improve the delivery of public services.²⁰ The overall program would be managed by one of the National Laboratories or by a newly created FFRDC with the mission to leverage technology to improve the delivery of public services. These testbed projects would be supported by local and state grants, cross-cutting federal government efforts, and public-private partnerships (PPPs) to employ AI to improve healthcare, workforce training, food production and distribution, and other areas. The overarching goals are to increase public trust in, understanding of, and confidence in AI; to learn how to use AI in ways that reduce inequality and enhance, rather than replace, human work; and to improve access, affordability, and availability of such services. At local, state, and federal levels, individual government agencies will gain long-term benefits by acquiring the necessary data infrastructure to employ AI to improve the delivery of public services.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Recommendation: Develop privacy-preserving technologies for the digital economy and demonstrate in a full-scale test their conformance with the General Data Protection Regulation.

The administration should direct NIST to establish and test privacy-preserving technologies that enable a risk- and outcomes-based approach to trust in the digital economy. The test should evaluate, at scale, conformance with relevant GDPR rules, conformance with existing US laws governing data privacy, and robustness with respect to innovations and advances in information technologies and data capabilities, especially those based on AI, machine learning, and the IoT. This work should include the development of technical definitions of privacy and application-specific measures of the utility of analyses that are based on privacy-protected data. The tests should include end user evaluations.

The administration should establish a near-term program that demonstrates privacy-preserving technologies to aid the trusted collection and sharing of data for the purpose of improving individuals' access to healthcare during large-scale biological events. This program should be jointly managed by NIST, the Department of Health and Human Services (HHS), the National Institutes of Health (NIH), and the National Science Foundation (NSF). This program will monitor system performance to inform the development of standards for the ethical use of the shared data and how data governance will be formulated.

²⁰ A potential source for the types of initiatives of interest is the OECD Network of Experts on AI (ONE AI). This group provides policy, technical and business expert input to inform OECD analysis and recommendations. "OECD Network of Experts on AI (ONE AI)," OECD.AI, accessed March 26, 2021, <https://www.oecd.ai/network-of-experts>.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Michael August

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Name: Michael August

Comments:

To address question 2, item D, I provide the following recommendation:

OpenAI Inc. (openai.com) is a research lab for Artificial Intelligence. Should the NAIRR be established, my recommendation is to establish a partnership with OpenAI Inc., in order to gain access to tools and technologies developed within OpenAI, including, for example, the Generative Pre-trained Transformer 3 (GPT-3) for natural language modeling. A public-private partnership could be established, for example, with either Microsoft, which has exclusive licensing rights to GPT-3, or directly with OpenAI Inc. for development of DoD-specific access rights for use of GPT-3. A public-private partnership with OpenAI Inc. could also be helpful for providing DoD employees with access to additional OpenAI tools and technologies as they are contributed and developed within OpenAI.

To address question 5, I provide the following recommendation:

The National Science Foundation runs an Industry-University Cooperative Research Center called the Cloud and Autonomic Computing Center, which could be used as an exemplar of an academic-industry-government partnership vehicle whereby resources can be shared amongst participating member institutions.

There are multiple universities, companies, and federal organizations which are members and participate in research and development activities, leveraging resources from other member institutions. This collaborative mechanism could provide federal organizations that are performing Artificial Intelligence research and development activities with access to world-class top researchers within Artificial Intelligence related fields from both academia and industry. In addition, this collaborative vehicle could also provide a commercialization mechanism by which academic Artificial Intelligence research products could be incubated by private sector members, thereby providing an incentive for engagement and active participation by all member institutions.

Federal Register Notice 86 FR 46278, <https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

BeeHero

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

September 30, 2021

VIA ELECTRONIC SUBMISSION

Re: Request for Information (RFI) on an Implementation Plan for a National Artificial

Intelligence Research Resource

BeeHero is a Fresno, California-based startup leveraging machine learning and low-cost sensors to increase crop yield and improve the health of pollinators.¹ We appreciate the opportunity to submit these responses, as a national artificial intelligence (AI) resource could be massively valuable to BeeHero and similar startups.

Pollination is a required and managed input in agriculture, just like fertilizer and water, and BeeHero is focused on optimizing pollination input through bee hive placement and delivery. Nearly three-quarters of the world's crops rely on bees for pollination, and while that used to come naturally, there are no longer enough bees in the right places to support pollination needs: bees face an increasing mortality rate, climate change is causing shifts in bee phenology and crop phenology, and the seasons and locations for crops and pollinators may no longer sync up.

BeeHero operates in a new, emerging field of data-driven pollination, enabling farmers and beekeepers to optimize pollination and increase crop yields. We collect and monitor data about hive health, weather, crop needs, landscape, etc., and combine that to estimate pollination input and how to deliver it. Our technology also allows farmers and beekeepers to monitor bee behaviors and visualize different patterns of stress placed on hives due to lack of food sources, viral disease, colony collapse disorder, or other factors. Building these models requires large amounts of data and skilled workers to engineer them.

The government can play an important role in the development of AI-enabled technologies to help solve the problems of tomorrow, which in BeeHero's case include solving problems of climate change and food security. The National Artificial Intelligence Research Resource

¹ BeeHero, <https://www.beehero.io/>.

(NAIRR) represents a key step toward promoting AI research and development by making available relevant data sets and by strengthening the talent pool through education and training.

1. What options should the Task Force consider for any of roadmap elements A through I above, and why?

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

The Task Force should prioritize providing data sets with cross-cutting relevance and import, connecting researchers, and growing the AI workforce.

First, a national AI resource could include data sets that would be directly relevant to BeeHero's work and to the work of many other researchers and innovators studying critical issues. In an area as complex as pollination, BeeHero studies a huge number of variables, and there is a massive range of other data we could use. And climate change is having a significant impact on crops, agriculture, and the variables relevant to our work. As such, inputs like satellite data showing vegetation and how it changes over time, aerial photography data, and weather trend data—the type that is likely already being collected—would allow us to better understand the state of agriculture and changes that impact pollination. If BeeHero were able to further integrate our bee data with those other data sets, it would enable us to predict agriculture needs and pollination yields with more accuracy. This sort of information is, in turn, vital to food security.

Second, and relatedly, a national AI resource can also serve a valuable convening function. As noted, climate change and modeling are huge areas for our work. But for a small company like BeeHero, modeling climate change cannot be our core business. If relevant data sets exist and modeling is happening elsewhere, connecting with that work would benefit us in substantial ways.

Indeed, given how pervasive the effects of climate change are, many small businesses could benefit from connecting to these same data sets and research. BeeHero already encounters those working on intersecting issues, for example companies leveraging AI to improve water management. The national AI resource can supply these innovators with common, valuable data sets and modeling, and be a place for us all to tap into parallel research. Connecting these businesses, researchers, and academics would advance work on numerous issues of national import.

Third, the Task Force should also focus on the core need for data scientists, addressing issues of education, training, and access to high-skilled talent. BeeHero's entire business depends on the quality and strength of our researchers, and as we aim to grow in the U.S. we will need to find

qualified talent to fill the new jobs we plan to create. AI promises to be huge, with an enormous number of potential applications, and this next industrial revolution will need people to fuel it. Likewise, for the U.S. to maintain a global lead in this emerging area, the country will need a constant stream of highly qualified people. The government should certainly consider investment in training and education. But broader investment in the AI field—from research, academics, and the private sector—will grow the field and the pool of talent that everyone will benefit from.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Please see response to question 1, above.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

The national AI resource should consider how to leverage existing government data sources and related efforts to better connect AI startups and innovators with the data they need. Predictive models, like those developed by BeeHero, require vast amounts of data, which must be acquired, stored, managed, and transformed into useful inputs. These tasks can present significant costs for startups with limited resources.² The government already has resources like data.gov that could help with this—data on topics including climate and agriculture that could be regularly updated and integrated into the national AI resource. The National Science Foundation and other federal agencies also award grants to fund data science research centers.³ Those grants and centers advance education and can produce useful data assets. The Task Force should build on this by coalescing data resources, services, and grant outputs in one place, and by facilitating their use in AI applications.

² Ivy Nguyen, *Could Data Costs Kill Your AI Startup?*, VentureBeat (Nov. 10, 2018), <https://venturebeat.com/2018/11/10/could-data-costs-kill-your-ai-startup>.

³ *E.g.*, National AI Research Institutes, <https://beta.nsf.gov/funding/opportunities/national-artificial-intelligence-research-institutes>; National Network of Big Data Regional Innovation Hubs, <https://beta.nsf.gov/funding/opportunities/big-data-regional-innovation-hubs-bd-hubs>; Institutes for Data-Intensive Research in Science and Engineering, <https://beta.nsf.gov/funding/opportunities/harnessing-data-revolution-institutes-data-intensive-research-science-and-0>.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Booz Allen Hamilton

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

INTRODUCTION

Booz Allen Hamilton (Booz Allen) is pleased to submit our response to the OSTP/NSF request for information (RFI) for the National AI Research Resource implementation plan (NAIRR). As the largest provider of artificial intelligence services for the Federal government¹, Booz Allen provides professional and technical services to research, design, architect, engineer and integrate AI solutions to accomplish critical missions and maintain U.S. technological leadership. We support some of our Nation's most high profile and innovative programs—including the Joint Artificial Intelligence Center (JAIC), Office of the National Coordinator for Health Information Technology, and the Defense Threat Reduction Agency (DTRA) Operations and Integration Directorate—to transform and advance their enterprise AI initiatives in a deliberate, outcome-focused manner to drive mission impact. More broadly, Booz Allen's AI business encompasses:

- An industry leading portfolio of AI/ML projects across civilian, intelligence and defense organizations ranging from early research to large-scale enterprise operations
- Award winning AI research and development teams publishing in top academic journals and forums²
- A Tech Scouting network and unique partnerships with big-tech AI/ML vendors and non-traditional start-ups (e.g., NVIDIA Consulting Partner of the Year 2018-2020; Databricks Federal Partner of Year 2021, AWS ML and MLOps competencies)

Our work recognizes that AI is not a single breakthrough technology, but a complex integration of people, processes, and technologies with a responsibility to use AI in a way that centers around people.

1.0 WHAT OPTIONS SHOULD THE TASK FORCE CONSIDER FOR ANY ROADMAP ELEMENTS ABOVE, AND WHY?

1.1 {Implementation Roadmap-A} Goals for establishment and sustainment of a NAIRR and metrics for success

Booz Allen suggests successful establishment and sustainment will revolve around three core attributes:

- **SHARED INFRASTRUCTURE:** Create, implement, and manage compute infrastructure resources that serve as a common sandbox for broader AI R&D and is representative of the changing operational environment with measurements to assess speed, progress, and sharing.
- **TOOLS, METHODS & DATA:** Make available a wide variety of rich data sets, world-class algorithms, and example uses cases that, when combined with the computing resources, will better democratize AI and allow for more contextual understanding and development. Measurements should include the amount and type of data sets (to include synthetic) and the amounts and types of openly available models, as well as the use/consumption of both the data and models.
- **EDUCATION & TRAINING:** Ensure AI-related assets are widely available to existing users while accommodating a potentially rapid growth rate; etc. This should include prioritizing education, training, and opportunities to engage repeatedly with real mission users to discuss the outcomes needed to achieve adoption. Measures should include the amounts and types of engagement and learning types.

¹ Bloomberg Government Market Analysis

² Thought Leadership for AI, Analytics, and Data Science (boozallen.com)

1.2 {Implementation Roadmap-C} A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources

We propose the creation of the NAIRR Council to coordinate between the organizations that hold a stake in the success of National AI research initiatives. This Council could follow the mold of other high-level coordinating organizations such as the National Security Council (NSC) and Joint AI Center (JAIC), and would facilitate working groups, establish strategic priorities, coordinate resources, and direct stakeholder efforts. This council can provide oversight to ensure progress against strategic goals and clear priorities. In this unprecedented partnership across industry, academia, and government, the council should be proactive, outcome-focused, and sensitive to feedback. They would allocate resources judiciously by providing and maintaining clear priorities. This will include leveraging Agile principles for adoption to consistently prioritize, groom, and refine the roadmap with end user feedback. As a part of the “Council” charter, they could define clear roles across decision-makers to best leverage strengths. Additionally, they could develop governance, IP protection policies and invest in researchers through the creation of public-private partnerships that enable technological innovation and collaboration by deploying best-practices and protecting data and technology³.

To build strong and collaborative public-private partnerships, we recommend: (1) **Accountability**: as partnerships evolve and collaborate, it’s important to have policies in place that hold the partnership accountable for successful collaboration⁴; (2) **Diversity**: collaborating with diverse companies and academic partners is a catalyst for innovation and rapid progress; and (3) **Integration**: provide a marketplace for AI and other strategic technologies to be implemented within the government through strategic investments in public-private partnerships to drive R&D⁵.

Recommended key functions: (1) Establish and review metrics and measurements of success; (2) Establish and review roadmap/implementation plan to ensure progress; (3) Continuously prioritize and update “backlog” based on feedback, user-testing, progress, and data; (4) Troubleshoot any problems that may arise hindering progress; and (5) Report to Congress and brief external stakeholders

1.3 {Implementation Roadmap-D} Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country...educational tools and services...and scalability of such infrastructure

In our recently published O’Reilly report, Enterprise AI Operations, we highlight the capabilities and components needed to successfully adapt and employ enterprise-wide artificial intelligence capabilities.⁶ A mature AI operations pipeline that will ensure scalable infrastructure and enable analytics and AI for NAIRR research needs should be thought through from the beginning to avoid rework and better support outcomes.

For more specific capability descriptions, see our response to Question 6. Below are some key overall areas associated with creating a framework to take capabilities from the lab to operations for NAIRR consideration.

- **Responsible AI Adoption** | Ensure AI solutions, when deployed, meet performance requirements, adhere to organizational standards and values, and are designed for adoption to achieve real mission outcomes.
- **AI Ops** | Integrate processes, strategies, and frameworks to operationalize AI and address real-world challenges and realize high-impact, enduring outcomes.

³ National Security Commission on AI Final Report (NSCAI), page 205

⁴ Casady, Carter, et.al., “A ‘New Governance’ Approach to Public-Private Partnerships: Lessons for the Public Sector,” 2017, Stanford University.”

⁵ National Security Commission on AI Final Report (NSCAI), pg. 449

⁶ Booz Allen Hamilton, “Enterprise AI Ops: A Framework for Enabling Artificial Intelligence,” O’Reilly Report, August 2021.

- **Data Engineering and Data Operations (DataOps)** | Locate required data and develop repeatable pipelines to increase value and make enterprise data accessible while promoting re-use.
- **ML Engineering and ML Operations (MLOps)** | Develop advanced algorithms using supervised, unsupervised, reinforcement, deep learning, etc. as required to support complex decision making.
- **Systems Engineering and DevSecOps** | Apply a structured framework for integration, documentation, and automation to develop, deploy, and monitor software and system solutions across an organization. Development, security, and operations (DevSecOps) integrate the critical components of security and focus on operationalizing applications through software and systems engineering.
- **Reliability Engineering** | Establish focused, clear objectives for AI solutions that are realistic with well-defined and quantifiable measures of success.
- **Infrastructure and Cybersecurity Engineering** | Protect data and AI applications for the long-term success of operationalizing AI with a strong technical architecture and cybersecurity policies.

1.4 {Implementation Roadmap-E} An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource

See our response to Question 6 below. We recommend looking at data, compute environment, and AI development holistically rather than independently.

2.0 WHICH CAPABILITIES AND SERVICES PROVIDED THROUGH THE NAIRR SHOULD BE PRIORITIZED?

One challenging aspect when establishing the NAIRR is appropriately bounding services. Put simply, NAIRR shouldn't try to become all things to all people. To establish and implement NAIRR at a rapid pace, we recommend focusing on three key areas: standardizing an AI Reference Architecture needed for a compute infrastructure, providing data sets and compute resources, and making available educational tools and resources.

2.1 AIOps—AI Reference Architecture for a Compute Infrastructure

To operationalize and deliver responsible, scalable AI solutions, Booz Allen developed an approach that starts with our recently released AI Reference Architecture (RA). A reference architecture provides technology agnostic guiding principles to standardize and accelerate the delivery of AI through common components, capabilities, processes, and terminology (Figure 1). Adopting a standard RA enables an organization to produce real-world solutions from an abstract framework, which drives consistency and standardization, surfacing risk early and supporting mitigation and reduction measures.

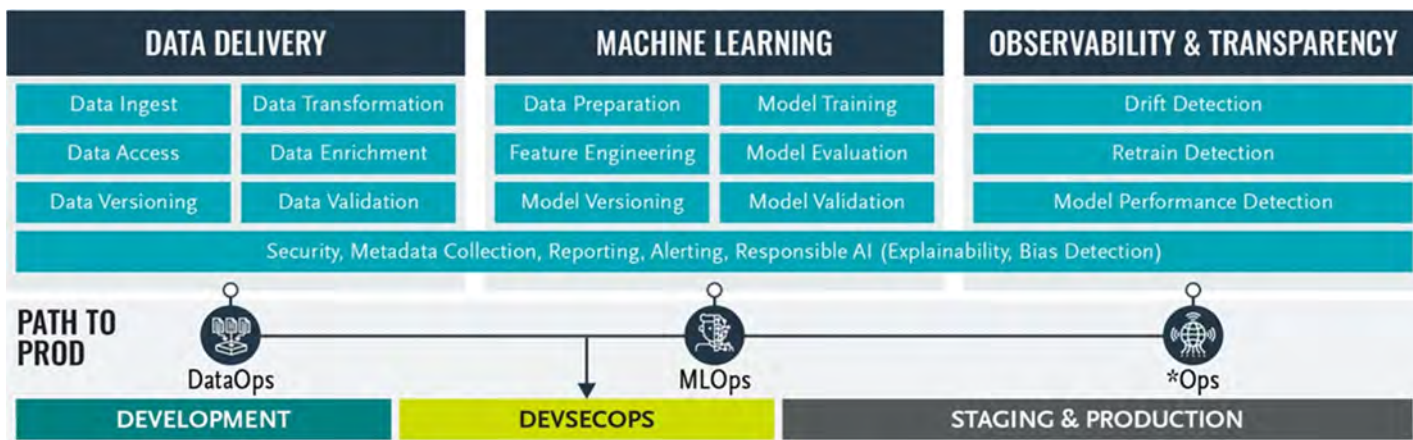


FIGURE 1: BOOZ ALLEN'S REFERENCE ARCHITECTURE (C) 2021 BOOZ ALLEN HAMILTON. ALL RIGHTS RESERVED.

- **DATA DELIVERY:** Covers activities that span ingestion, storage, transformation, enrichment, and delivery of data for analytics at scale. Robust, repeatable data delivery provides a critical foundation for analytics activities, and provides lineage and provenance collection to support data governance.
- **MACHINE LEARNING:** Encompasses the performed activities to prepare and transform data into insights and inferences that serve client and other business needs. The RA defines the main capabilities that when combined comprise a holistic ML Workflow.
- **OBSERVABILITY AND TRACEABILITY:** Provides processes and capabilities for monitoring and reacting to changes in ML workflow performance. This includes tracking performance, detecting when performance shifts beyond acceptable limits, and triggering appropriate actions in response.
- **CROSSCUTTING COMPONENTS:** Provides core functions including security, metadata collection and explainability to support the development of analytic services.
- **PATH TO PRODUCTION:** Combines software delivery best practices with Data Operations (DataOps), Machine Learning Operations (MLOps) to establish a robust path to production.

2.2 GOVERNMENT CURATED DATA SETS AND COMPUTE RESOURCES {ITEM D}

Data and computational resources are foundational to successfully creating responsible AI; therefore, related capabilities and services should have priority. This should include data, curation, hosting, auto tagging, maintenance of a variety of types of data sets for different types of use cases that have been cyber secured. Additionally, we recommend incentivizing the sharing of data sets for research purposes. We would also suggest collaboration with the Office of the National Coordination for Health Information Technology (ONC) on their recent report "[Training Data for Machine Learning \(ML\) to Enhance Patient-Centered Outcomes Research \(PCOR\) Data Infrastructure](#)"⁷ as the main goal for this effort is to make data sets available for AI/ML research. Additional information is in our response to Question 1.4 Implementation Roadmap D above and Question 6.

2.3 EDUCATIONAL TOOLS AND RESOURCES {Item D}

Getting to adoption is more than just providing data and computing resources. It's also creating a culture of learning and sharing, which includes datasets, models and other resources that emphasize an open-source mentality. Prioritizing retraining and collaboration with educational institutions will help integrate AI curriculums into everyday learning, starting from a young age with reinforcement as educational groups move through their learning journey. It's crucial to remove the barrier to entry for learning about and engaging with AI to shape the workforce required for the United States to stay at the forefront of AI innovation. For example, Booz Allen has invested in "The AI Education Project" to help lower barriers to AI education. The project works with K-12 schools to increase access to computer science education to help students build a foundational knowledge of how AI works and how it affects their lives⁸. Introducing AI early in a student's education helps remove AI's mystique and further integrates it into society as a fundamental tool. Partnering with a range of higher-education institutions will provide high-quality technology education to foster a creative, innovative, and AI-fluent workforce.⁹ The University of Florida entered a public-private partnership with NVIDIA to build an AI-centric data center where students receive hands on experience with AI and industry leading tools like

⁷ [The Office of the National Coordinator for Health Information Technology- "Training Data for Machine Learning \(ML\) to Enhance Patient-Centered Outcomes Research \(PCOR\) Data Infrastructure" Report \(healthit.gov\)](#)

⁸ <https://medium.com/the-ai-education-project/introducing-the-ai-education-project-3c1f1fc31fd2>

⁹ [National Security Commission on AI Final Report \(NSCAI\)](#), p. 543

supercomputers¹⁰, advancing research efforts. Exposure to environments like a large, shared infrastructure, AI research programs, and curriculum will help solidify U.S. colleges as technology innovation hotbeds.

3.0 HOW CAN THE NAIRR AND ITS COMPONENTS REINFORCE PRINCIPLES OF ETHICAL AND RESPONSIBLE RESEARCH AND DEVELOPMENT OF AI, SUCH AS THOSE CONCERNING ISSUES OF RACIAL AND GENDER EQUITY, FAIRNESS, BIAS, CIVIL RIGHTS, TRANSPARENCY, AND ACCOUNTABILITY?

Responsible AI involves more than just fairness and ethical outcomes. These concepts are important, but they only represent a system's characteristics—which are neither good nor bad on their own. As a result, assessing whether an AI system meets responsible and ethical requirements requires two key elements: 1) clearly articulating what responsible AI development looks like and adopting from inception, and 2) incorporating risk mitigation strategies to identify, assess, and address possible bias.

3.1 CLEARLY ARTICULATE RESPONSIBLE AI DEVELOPMENT

Thinking about responsible AI from the design phase all the way through to implementation and monitoring is critical to ensure it represents ethical outcomes.¹¹ Responsible development is a core aspect that needs to be at the forefront—focusing on the combination of elements that make up responsible AI, such as adoption, ethics, and the workforce.

Responsible AI becomes meaningful only when considering how an artificial intelligence system impacts people. The “*who*” question is critical. In other words, responsible AI needs a focus—whether that's an individual, a group or class of people, or an entire society. But responsible AI *also* needs a source of authority, or moral compass. That can be an organization's shared values and principles, or a society's norms and cultural practices. It is only then that we can address, for example, what a “*fair*” AI system is, for whom, and why we consider the implications to be fair in the first place. Embedding these values into the governance structure with the proper controls and measures to validate execution is key. This governance structure, including controls and metrics, should have agreement at an organizational level and not on a per project basis.

There are five main considerations when integrating Responsible AI:

1. Understand how artificial intelligence systems will impact an organization's stakeholders in specific and tangible ways. This assessment should include routinely considering the political, economic, social, technological, legal, and environmental impacts across different stakeholders over time.
2. Build meaningfully diverse and inclusive development teams. Include members with different backgrounds, skills, and thinking styles. A team's collective experience and insights will reduce unconscious bias, identify potential unintended consequences, and better reflect stakeholders' wide-ranging values and concerns.
3. Develop mechanisms for data provenance and auditability to verify AI systems are operating as intended. If something goes wrong, data tracing and auditability mechanisms will help uncover data or concept drift or potentially expose upstream/downstream data issues. Clear accountability mechanisms and data test can help reduce ethical concerns (e.g., data bias and amplification of the bias in ML for training, inference, etc.), so it is critical to transparently account for the results. Teams should understand that they are accountable for the actions, outputs, and impact of their models.

¹⁰ University of Florida, [UF Announces \\$70 million artificial intelligence partnership with NVIDIA](#), v July 2020.

¹¹ Booz Allen Hamilton, [“Enterprise AIOps: A Framework for Enabling Artificial Intelligence,” O'Reilly Report, August 2021](#)

4. Stay informed regarding AI technical developments. Because this field of study changes rapidly, tools used to design and implement ethical systems have limited shelf-lives. A model's sophistication will often outpace ethical tools, increasing the probability something will go wrong and reducing the ability to fix it if it does. Maintaining awareness of AI technical developments and implementing measures to monitor can help mitigate risk and protect an organization from unintended consequences.
5. Design systems with specific applications and use cases in mind. Assessing the "fairness" of a model requires *context* and *specificity*. AI systems should be fair, but fair to *whom*? And in *what way*? Fairness is a laudable goal but only becomes useful when applied to a specific situation. Something that may be a fair outcome for someone in one situation could appear totally unfair in another situation.

3.2 INCORPORATE RISK MITIGATION STRATEGIES TO IDENTIFY, ASSESS, AND REDRESS POSSIBLE BIAS

As described in our recent “[NIST Artificial Intelligence Risk Management Framework Request for Information \(86 FR 40810\)](#)” response, we are deeply committed to maintaining an acute awareness of the societal and environmental impacts of our AI systems and applications, and we ensure that our company and our people design AI systems that are grounded in real-world implications. To provide checks and balances and ensure the implementation of AI guiding principles, we believe that the implementation of an AI governance process, like that described in Figure 2, is necessary. In addition to ensuring the evaluation of AI projects to systematically mitigate risk, governance builds stakeholder confidence in AI through an organizations’ responsible use of AI. Well-designed controls promote the application of AI through effective management, ensuring that it is meeting performance requirements and ethically used.

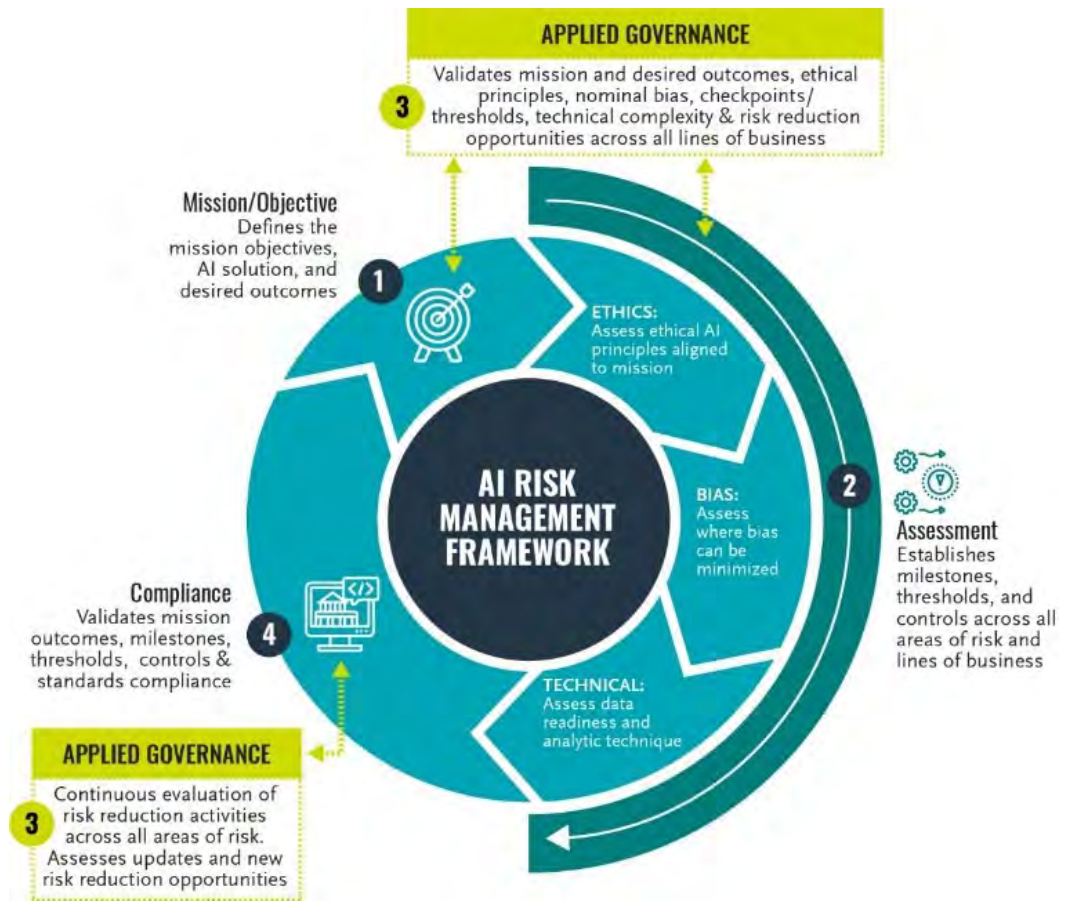


FIGURE 2: BOOZ ALLEN AI RISK MANAGEMENT FRAMEWORK (C) 2021 BOOZ ALLEN HAMILTON. ALL RIGHTS RESERVED"

4.0 WHAT BUILDING BLOCKS ALREADY EXIST FOR THE NAIRR, IN TERMS OF GOVERNMENT, ACADEMIC, OR PRIVATE-SECTOR ACTIVITIES, RESOURCES, AND SERVICES?

The private sector is rapidly pushing for AI growth through forming new partnerships between technology powerhouses and investing in smaller firms and start-ups, creating scalable multi cloud environments,

developing a range of AI frameworks for rapidly developing modeling tools, and creating/implementing large scale workforce programs. Booz Allen is committed to investing in AI innovation, including investing in non-traditional businesses with potential to positively disrupt public sector missions. For example, we invested in Latent AI with the goal to leverage the partnership to help clients implement AI models into sought-after end user devices to help drive AI adoption.¹² Investments and collaboration between private-sector entities foster rapid innovation and implementation of technology at a greater scale.

Workforce training will be one of NAIRR’s most powerful tools. Many private-sector companies have established and well-tested training programs to educate, train, and grow their workforce. Through a partnership with NIVIDA’s Deep Learning Institute (DLI), Booz Allen has developed a Deep Learning Training service to help clients build data science skills required, to become experts that can implement Machine Learning technology within their company or agency.¹³ This type of customized facilitated training service will help accelerate analysis and drive mission success.

There are numerous building blocks that already exist across the Federal Government, Academia, and industry. We would suggest the Government conduct a larger data call to get more information about those efforts. To be most effective, we suggest setting up some overall categories/bounding boxes with clear definitions so that the collective group can better organize their inputs in a digestible manner. New efforts arise every day and would encourage non-traditional thinking of how-to best leverage and learn from those efforts. With that being said, we offer a few specific building blocks below for consideration in addition to the other efforts listed throughout our response:

NAME	BRIEF DESCRIPTION
DoD ADVANA Data Platform	A central hub for advanced analytics, data science and AI connecting Senior Leadership to authoritative data sources needed for data-driven decisions. ¹⁴
VA/DOE- VA’s Million Veteran Program	The Department of Veterans Affairs (VA) and the Department of Energy (DOE) are partnering to drive technology innovation and transform health care delivery for Veterans. The partnership brings together VA’s healthcare and genomic data with DOE’s high-performance computing (HPC), artificial intelligence and data analytics. ¹⁵
The AI Education Project	The AI Education Project is a 501(c)(3) non-profit centering equity and accessibility in AI education. They educate students, especially those disproportionately impacted by AI and automation, with the conceptual knowledge and skills they need to thrive as future workers, creators, consumers, and citizens. ¹⁶
DoD GameChanger Open Sourcing	Over 15 thousand documents govern how the Department of Defense (DoD) operates. These documents exist in different repositories, often on different networks, are discoverable to different communities, updated independently, and evolve rapidly. GAMECHANGER offers a scalable solution with an authoritative corpus comprising a single trusted repository of all statutory and policy driven requirements based on Artificial-Intelligence (AI) enabled technologies. ¹⁷

¹² Booz Allen Hamilton, “Investment builds on AI capabilities supporting algorithmic warfare at the edge,” July 2021.

¹³ Booz Allen Hamilton, “Deep Learning Training”

¹⁴ [Be Ready for What's Next - Government Technology Insider](#)

¹⁵ [DOE and VA Team Up to Improve Healthcare for Veterans | Department of Energy](#)

¹⁶ [The AI Education Project](#)

¹⁷ [GitHub - dod-advana/gamechanger: GAMECHANGER aspires to be the Department’s trusted solution for evidence-based, data-driven decision-making across the universe of DoD requirements](#)

Data.gov	The U.S. Government’s open data site—access to data, tools, and resources to conduct research, develop web and mobile applications, and design data visualizations. ¹⁸
NIH NLM Data Science Training	NIH National Libraries of Medicine (NLM) Data Science Training Program that provided targeted training to all NLM’s 1,700 staff members. With a focus on becoming more aware of data science and its incorporation into so many NLM products and services. ¹⁹

6.0 WHERE DO YOU SEE LIMITATIONS IN THE ABILITY OF THE NAIRR TO DEMOCRATIZE ACCESS TO AI R&D? AND HOW COULD THESE LIMITATIONS BE OVERCOME?

The quality of resources and investments will have the greatest impact on NAIRR’s goal to democratize access to AI R&D. A holistic approach that better connects research to operations will allow for integrating AI R&D into the fabric of America. We encourage the Government to think of both basic and applied research and how to translate and integrate into the operational environment that exists today including the use of “operational test beds.” Below are the key areas that we see as critical for democratization, and some (not exhaustive) associated sub areas for consideration.

DATA QUALITY & AVAILABILITY | There is a need for broad and consistent access to high-quality operationally relevant catalogued data for training and testing of AI models. This should include assessing the ability to access data required, manage its quantity, and evaluate its quality, as well as, evaluating the capacity to understand the data (e.g.) data dictionary. Additionally, there is need for the ability to correctly evaluate biased data and make the proper corrections.

- **Data Readiness & Management:** Managing data as an asset can help NAIRR address these challenges and support AI needs through a multilayered approach to data readiness. A holistic approach includes the following key areas: (1) Data strategy: aligns data capabilities with AI needs and establishes effective data management and data usage practices tailored to AI needs, (2) DataOps processes optimize workflows for automated data ingest, processing, and storage to help keep up with data needs at scale, (3) Standardized methods enable data ingest, transformation, validation, and preparation, (4) Data governance expedites value and spans compliance, risk, and regulation related to data (including privacy, security, and access controls), and (5) Responsible data policies specify access rights, “right to see” authorizations, ethical principles, and acceptable applications for data usage across the organization.
- **Synthetic Data:** There is not enough diverse, high-quality, labelled data accessible, thus limiting the application of trained ML models to many domains. Synthetic data is the generation of artificial data with the aim of reproducing the statistical properties of an original dataset generated by real world events. This data can be either partially or fully synthetic, with the former containing generated data alongside original data, and the latter composed exclusively of generated data. It is a powerful solution for those who do not have the resources to collect, label and process huge amounts of data typically needed to train complex algorithms. The Department of Veteran Affairs has been actively using synthetic data through its [Veterans](#)

¹⁸ [Data.gov](#)

¹⁹ NIH- [Building Data Science Expertise at NLM – NLM Musings from the Mezzanine \(nih.gov\)](#)

Health Administration Innovation Ecosystem (VHA IE) because of restrictions associated with PHI and PII.²⁰

- **Open Source:** Maximize data access by streamlining policies and regulations to promote and increase data sharing across organizations. Supporting and investing in AI tools like JAIC's GameChanger²¹, which is modernizing data access rules, will ensure more users have access to relevant and high-quality data. Other evolving data sharing practices include multi environment designs based on security where verified researchers could have access to more sensitive information.²² Incentivizing the private sector and academia to share data sets is a critical element for democratizing AI R&D.

GOVERNANCE | There is a need to ensure the creation/use of “Responsible AI” and risk management frameworks that include oversight and sustainment for all data, model, and computing environment efforts focusing on reusability. This includes processes for configuration management, testing, and verification & validation.

- **Responsible AI- Risk Management:** See Question 3 above and extensive response to the NIST Artificial Intelligence Risk Management Framework Request for Information (86 FR 40810)²³
- **Data Test & Evaluation:** Integrated data tagging and labeling capabilities while also providing the ability to develop ground-truth datasets.
- **Model Management:** As with any other form of software, ML models need documentation, testing, and version-control to help the team maintain knowledge of its structure and functionality. Unlike conventional analytics, we must rely on surrounding documentation, testing, and metadata.

ARCHITECTURE & DESIGN | There is a need for a modular software architecture that deployable in a timely manner while ensuring a focus on security and scalability. It needs to include the use of open, secure software architectures comprised of best-in-class COTs, GOTs, and open-source components that balance quick experimentation with affordable scalability. The system should support interoperability and be reliability operated and maintained.

- **Evolutionary data architecture:** Data architectures designed with flexibility and adaptability in mind to evolve at the rapid pace of innovation. Maintaining the proper design abstraction allows for the component replaceability necessary to keep up with AI’s rapid evolution.
- **Scalable and Flexible data pipelines architecture:** Designing data pipelines for scale and flexibility at the beginning for efficiency rather than attempting to scale at subsequent stages of growth.
- **Institute open API interfaces for data sharing:** An application programming interface (API) defines how to push or pull data during a data exchange. APIs help bridge data silos and break them down into more usable parts that greatly increase integration and reduce development time.

SECURITY | There is a need for broad and consistent access to configurable, scalable, accredited computing environment to support development, experimentation, and operational hosting of AI capabilities. This requires ensuring systems built are robust and secure and incorporate security practices at every stage of development. Additionally, there should be methods to properly monitor new complex AI solutions and have techniques in place to combat adversarial AI.

²⁰ How synthetic data will improve Veteran health and care - VAantage Point

²¹ National Security Commission on AI Final Report (NSCAI) pg. 305

²² National Security Commission on AI Final Report (NSCAI) pg. 449

²³ <https://www.regulations.gov/comment/NIST-2021-0004-0058>

- **Adversarial**- As more and more systems integrate AI, there is growing concern about how to protect these systems from attacks (e.g., evasion, poisoning, model-stealing, backdoor, etc.). This requires a novel approach to AI protections using methods from the well-established domain of control theory to provide mathematically provable protections built into the system during development. This provides engineers with a flexible and generalizable way to design multi-layer neural networks that are resilient against many attacks without compromising the system’s ability to learn.
- **Securing the Infrastructure**: Application and mission owners must architect their cloud environments to not only be compliant with security and risk standards but simultaneously ensure the implementation of the right cybersecurity designs and tools to protect against emerging threats and attacks. This requires continued enhancements as the platform becomes operationalized, resulting in the need for organizations to migrate to a SecDevOps continuous development culture.

CULTURAL & ADOPTION | There is a need for a clearly articulated AI vision and strategy that directly supports the organization’s mission. This should include proper training, communication plans to support building of trust and transparency, and operational user involvement and feedback.

- **Education & Training**: There is not enough AI-fluent talent to meet the government's needs and a rapidly evolving technology landscape. Workforce training and educational tools are key to supporting a strong computing infrastructure. A focus on training so that all users are AI knowledgeable also provides access to a larger talent pool, to build meaningfully diverse and inclusive design and development teams. This should include educational and training courses ranging from intro to refresh course at varying degrees of technical competency based on roles individuals will play focusing heavily on achievable mission outcomes.
- **Mission Acceleration & Mission First Engineering**: To truly bridge the “compute divide,” the focus needs to also include the connection to the mission so that the collective community (research and mission) can learn and grow together. Many are too focused on prototypes and pilots that when brought to operations may never actually deploy as the operational environment was not part of the original design. The creation of AI is most effective when starting with the outcomes desired and fully understanding the “mission” needs. An AI team should be a cross-functional, integrated team working on holistic AI/ML solutions that leverage the expertise of the members (technical, operational, enabling) with the key understanding that all functional areas are critical for success. Operational feedback loop(s) and learning is one of the most essential pieces of getting AI into operations. Your teams must be able to monitor data, models, applications, and processes to evaluate potential changes/updates needed throughout their life cycle to ensure you are achieving the right outcomes responsibly. Like sports, we need to practice like we will play so connections to operational environments from both a testbed environment and connection to operational users is critical. Keeping research separate will continue to exacerbate the divide between getting to operational AI and real adoption with mission outcomes.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Cadence

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

RFI Public Response For:

National Artificial Intelligence Research Resource – NAIRR

Response: National AI Research Resource” in the subject line of the message. Attn: Wendy Wigen, NCO,
2415 Eisenhower Avenue, Alexandria, VA 22314, USA.

Bottom Line Up Front

“People who are really serious about software should make their own hardware” - Alan Kay

Artificial Intelligence curricula is available worldwide and on the internet. Where there is AI education there are compute resources. Though those resources may need to be bolstered, it is not a comparatively critical problem.

There is a broad and growing open source software community on the internet that provides foundational libraries for AI software development. These resources provide a substantial head start for anyone in the world wishing to engage in AI research or product development.

It is widely projected that a majority of AI systems and products will operate at the network edge, if connected at all. Latency, power and cost will be at a premium for competitive capability.

While AI software development has an organically growing worldwide foundation, mission optimized AI hardware investment is significantly lacking. This is the area where a protected national investment could create a strategic advantage for the country.

Fabrication of advanced microelectronics is expensive, especially compared to software. However, the payback in terms of capability and value versus software on standard hardware is more than worth it.

Fortunately, there are proven alternatives to taking every new custom hardware solution through physical realization for evaluation. Hardware accurate digital twins are a by product of the commercial best practices in microelectronics design. This type of digital twin can provide highly accurate analysis to project manufactured device characteristics. It is the link between electronic design automation (EDA) tools and the semiconductor manufacturers via process design kits (PDKs) that creates the accuracy. This link is the basis on which the entire commercial fabless industry has been based.

Cadence, the world leader in hardware accurate digital twins, would be happy to provide briefings to explain and discuss the technology.

Cadence is the only EDA vendor with trusted supplier accreditation.

Cadence has worked with the Air Force Research Lab (AFRL) and the Defense Micro-Electronics Activity (DMEA). Those organizations have been accumulating capability and expertise in commercial best practices and hardware accurate digital twins. They would be excellent candidates to act as the initial administrators and governors of a NAIRR for custom hardware.

Making government IP access rights a condition of the access to the NAIRR would ensure that the country would be a long-term beneficiary of this investment.

Google: AI education resources = > About 862,000,000 results (0.54 seconds)

As outlined in § 5106(b) of Public Law 116–283, the implementation roadmap developed by the Task Force should include the following:

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success

Goal of program, in part, should be establishment of an innovation pipeline that feeds our national defense needs. Core success metric should be deployed national security capability with, at least in part, their origin in the NAIRR.

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including: i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

Ownership for the NAIRR and its accumulated IP and IP access rights should be within the DoD. The Air Force Research Lab (AFRL) and the Defense Micro-Electronics Activity (DMEA). Those organizations have been accumulating capability and expertise in commercial best practices and hardware accurate digital twins. They would be excellent candidates to act as the initial administrators and governors of a NAIRR for custom hardware.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources

While national security priorities should be foremost, other major government departments with competent scientific staff should be a part of a multi-department governing board.

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

Shared compute resources are not an essential element of the NAIRR. Every major educational and research center already is equipped with compute resources. For many organizations that claim a paucity, it is a matter of prioritization.

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of highquality government data sets as part of the National Artificial Intelligence Research Resource

There is an established multi-branch hardware emulation center at AFRL. The commercial electronics industry has shown how such a capability can be successfully scaled. Data sets and IP need to go through a rigorous screening process (with limitations of that process know and noted). The DoD should be granted full access rights to all of the data and IP registered.

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls

There should be at least a two-level security scheme for incoming data (open and ITAR). It seems advisable to instill provisions for higher levels of classification and screening that leverage the same infrastructure investment.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research

N/A

H. A plan for sustaining the NAIRR, including through Federal funding and partnerships with the private sector

Over time, the savings to the DoD in overall system development and lifetime cost should be worth billions of dollars. A percentage of those savings could be applied to the sustainment (and expansion costs). A royalty scheme for commercially successful products whose basis IP was developed in the NAIRR could be an additive element (need caution here to not squelch business viability).

I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

Would encourage an initial structure and governance based on on-going infrastructure development within the DoD, then evolve as viability and path forward becomes clear.

1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

CalypsoAI Corp.

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

CALYPSO AI

CalypsoAI's Response to the Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource

Calypso AI Corp.
2955 Campus Dr. Ste 110
San Mateo, CA 94403

Under Congresswoman Anna G. Eshoo, CalypsoAI is an [official supporter](#) of the [bill](#) that created the National Artificial Intelligence Research Resource (NAIRR) Task Force. We are pleased to respond to the National Science Foundation (NSF) and the White House Office of Science and Technology Policy (OSTP)'s request for information (RFI) on an Implementation Plan for a NAIRR. CalypsoAI believes an initial roadmap created by the NAIRR Task Force and informed by academia, industry, and government, is an important step towards achieving the goal of democratizing access to a cyberinfrastructure that fuels Artificial Intelligence (AI) research and development (R&D). CalypsoAI's CEO, Neil Serebryany, highlighted the importance of this initiative in a [NAIRR Task Force Act press release](#) stating, "the research infrastructure that will be created by this legislation is critical to our nation's ability to lead the world in building secure and operational AI."

As a software company and thought leader for Secure and Trusted AI, CalypsoAI knows firsthand the importance of access to a shared holistic advanced computing and data ecosystem. Our work on the testing, evaluation, verification and validation (TEVV) of AI and Machine Learning (ML) models resulted in contracts with the U.S. Department of Homeland Security (DHS) Science and Technology (S&T) Office's Screening at Speed Program (SaS) and the Secretary of the Air Force Concepts, Development, and Management Office (SAF/CDM), named Secure Artificial Intelligence. We are also working tangentially with the National Institute of Standards and Technology (NIST) on Software for Trusted Intelligence.

Drawing from our cutting-edge AI R&D team and experience working with industry and government in the AI space, CalypsoAI provides the following comments to inform the Task Force's consideration of options and development of a NAIRR implementation roadmap.

RFI-Specific Question Answers:

1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list) for which you are identifying options.]

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

Planning frameworks are frequently adopted out of disciplinary habit and as a feature of cultural production within mature organizations. Just as software engineering will default to agile and iterative methodologies, or military contractors may rely upon the conventions of program management professional trainings (PMPs), national technology policy and planning frameworks typically reflect the beliefs, bias, and organizational structures of the governing body. The challenge, and opportunity, of national AI planning initiatives is to make a clean break, to reflect upon the diversity of planning frameworks, and to construct a planning methodology best attuned to the nature of AI problems – not configured according to the composition of stakeholders.

In this spirit, it is *too early* to settle on key performance or success metrics, or even to articulate the goals, as the problems are still undefined. However, we do know a few things; first, at this point in the nascent 90-year history of artificial intelligence and cybernetics, many core problems include the threatening reproduction, extrapolation, and distribution of misinformation/disinformation, all of which pose a threat to our national security. Second, we know that technology companies are ill-equipped to manage the threats and consequences of their work. Finally, extensive peer review research in the realms of sociology, psychology, and human computer interaction has taught us that many tools, products, and algorithms are deeply divisive and harmful to human relationships, good governance, and social welfare.

Recognizing our limitations, goals for AI resource planning should begin with a review on the emergent socio-technological failures of our AI era. Exploratory multi-stakeholder exercises in scenario planning, speculative design, advocacy planning, and dialogue mapping should precede conventional structured planning and roadmaps. Resource planning, strategy, and investment planning should then align to the composition of the problems that emerge - *not* the whims, expectations, or beliefs of the participants.

Since the research, problems, and key actors shift rapidly within the domain of AI with the emergence of each new technology, user trend, or novel application, metrics should not be established at an overarching policy level. The formulation of metrics should be relevant to the specific bodies that are then tasked to pursue each problem. However, NAI RR should be the body to coordinate oversight or the review of employed metrics, ensuring their implementation, relevancy amidst the rapid shifts within the AI space, adherence to best practices and policy, and consistency with metrics used across the government.

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

Governance should be widely decentralized across multiple existing bodies but should coordinate efforts to reflect the same sound principles, such as existing federal

memorandums on Responsible AI (RAI). While entities such as OSTP should provide oversight of the Research Resource by leading the conversation on AI regulation, each U.S. government agency should manage its own adoption, implementation, and decision-making. This approach is consistent with other large-scale strategic federal compliance concerns, such as information security and management. Existing bodies such as the Government Accountability Office (GAO) and U.S. General Services Administration (GSA) are already effectively positioned to support the release and support the implementation of standards.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

To points B and C of the initial implementation roadmap, NAIRR could refer to the Federal Lab Consortium (FLC) as a governing structure example of ownership, administration, governance, and oversight. It is governed by an Executive Board with four nationally elected positions, six regional coordinators, six Members-at-Large, and chairs of standing committees. This board determines policy, direction, and budget. Furthermore, this Executive Board is advised by the National Advisory Council (NAC) informed by user communities (industry, academia, government, federal laboratories).¹

For the NAIRR, these executive board members, regional coordinators, Members-at-Large, and chairs of standing committees should be nominated via an open call for nominations, similar to the U.S. Department of Commerce's current call for nominations for its National AI Advisory Committee. This provides members of government, industry, and academia opportunities to promote the best and brightest in their fields and to provide representation of various stakeholder entities. After nomination, a selection committee from NSF and OSTP should review nominations and select the NAIRR governing body members based on established records of distinguished service and eminence in the field. These members should serve for two-year terms in order to provide adequate time to execute initiatives while still accounting for the need to bring in new talent within such a fast-paced field.

This hierarchical structure of governance will be important for NAIRR stakeholders because it establishes authority and accountability. Creating a single entity like the FLC also helps accomplish the NAIRR's goal of democratizing access to the cyberinfrastructure that fuels AI research and development by establishing buy-in and visibility of stakeholders across government, industry, and academia.

However, if this body is to go beyond strategic planning and participate in programmatic decisions, it should do this in direct coordination and co-design with the federal agency, or agencies, with equities.

¹ Federal Lab Consortium, <https://federallabs.org/about-us/organization>

- D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

Initiatives such as Data.Gov or USA.Gov, initially developed by GSA Technology Transformation Service (TTS), demonstrate the successful instantiation of shared public infrastructures, and should be both replicated and scaled to create a shared computing infrastructure. With both examples, multiple government agencies participate in managing computational resources with high quality user experience, upholding relevant standards, establishing modern development practices, and responsibly implementing open-source licensing. While it does not matter which agency leads the creation of a similar initiative, it is imperative to establish rules and standards for how information is stored and shared. This consistent curation, coupled with a simple user-interface portal, will enable researchers across the country and various industries to better use the information we have. It makes the information more usable, both across data sets and as a way of tracking over time, which promotes good governance and auditability. Without one particular agency leading the initiative for setting standards on the storing and sharing of information, every agency would set their own standards, losing the power of the NAIRR's shared infrastructure.

- E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

Strong precedent exists for the federal management of high-quality data. Over the last ten years, agencies such as the U.S. Department of the Interior, the National Aeronautics and Space Agency (NASA), the U.S. Department of Health and Human Services (HHS), and the Environmental Protection Agency (EPA) have all taken active steps to make computable quality data available to the public. Yet solutions such as Data.Gov have been more effective because the intentional design – such as the use of Restful APIs, GeoJSON, robust metadata and similar normative data standards – have enabled rapid adoption and consistent updates. Rather than create new data resources, the NAIRR should invest in and scale the existing proven solutions, such as Data.Gov and USA.gov, to best align with the needs of the AI research and commercial communities.

A large barrier to the dissemination of high-quality data for the NAIRR is the quick and constant evolvment of metrics within the AI space. Without an agile approach to a data repository, data could already be outdated by the time it gets into the hands of the NAIRR stakeholders. To mitigate this problem, agencies should be required to input data within a certain timeframe after acquiring it. Additionally, the NAIRR should advocate for individual data efforts --including those of NASA, HHS, and EPA --to be consolidated into Data.Gov. This will provide the NAIRR stakeholders with one go-to database for high-quality government data sets that have breadth and depth of the most cutting-edge information. Finally, the storage and sharing of the data within such a data repository

should follow the standards set by the particular agency chosen to create and implement the standards, as discussed above.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

Typically, privacy and civil rights and civil liberties requirements are met through large reporting instruments. Although this approach is comprehensive, it risks infringement because it is not monitored consistently. Consequently, rather than implementing large reporting instruments, privacy and civil rights assessments should be more consistent and more modular in composition. This includes implementing prioritized checklists in place of sprawling compliance documents, constant discussion across stakeholders, and quickly collaborating when an issue arises to ensure it does not spread or recur. By aligning these privacy and civil rights and civil liberties checks with the pace of AI development, we may better protect citizens and build their confidence in AI systems.

H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; and

AI governance suffers from extensive thought leadership and minimal real-world experimentation. There are few examples of companies or universities leading exploration in everyday human environments – such as the construction of an experimental intersection at the University of Michigan for autonomous vehicle testing – on the implementation of concepts in AI ethics or design principles. Rather than continuing to allocate money into the publication of white papers, reports, and large documents that explore ideas, the NAIRR should be sustained through action. This means that federal funding and partnership investment should support creative physical exploration and experimentation of AI governance concepts. For example, it is myopic to invest funding or pursue partnerships only with machine learning engineers, given that their work sits within human environments constructed and informed by a range of other professions. Partnerships should be broad in spectrum and interdisciplinary though targeting finite goals.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

The NAIRR, in promoting and democratizing access to the infrastructure needed to fuel AI research and development, must address the global gap between AI development investment and AI deployment. While billions of dollars were spent on AI in 2021, AI deployment is remarkably low. Therefore, the NAIRR should specifically prioritize research and investment in capabilities and services that address the security and risk management of algorithms. This will promote trust in AI/ML systems and speed up their development and deployment. This prioritization by the NAIRR will strengthen the United States' national security and maintain its global lead in AI by putting AI/ML security at the forefront of research, development, and deployment of AI and AI-enabled solutions across U.S. industry, government, and academia. As an industry partner in the AI/ML

TEVV and risk management space, CalypsoAI has seen firsthand the importance of this prioritization of AI security and holds a suite of tools to assist NAIRR and its stakeholders in this space.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

AI/ML products developed with minimal standard tooling leads to algorithms of uncertain quality, subjective trustworthiness, and potential vulnerability to attack. Across all sectors, organizations lack the tools to adequately assess and monitor AI/ML products, often leaving such solutions undeployed.

Built-in oversight tools and mechanisms to support progress and accountability such as periodic evaluation, design choices that enable metrics collection, and transparent reporting will be necessary to reinforce principles of ethical and responsible AI research and development. Specifically, with CalypsoAI's expertise in TEVV of AI/ML, the Task Force's roadmap must provide solutions that enable the following for the development and deployment of AI/ML models:

- I. *Responsible Application*: Provide rapid algorithm testing to encourage transparency and accountability in algorithm development and deployment, while respecting the intellectual property rights of participating vendors. Testing should also place the least possible amount of burden upon the vendor.
- II. *Accessible Results*: Can be easily accessed and used by sophisticated experts who are not data scientists specialized in model testing. The testing process must be swift to expedite the role of model review within the acquisition process. Test results should be concise, easy to understand, and contextualized so that model review teams can rapidly ascertain, compare, and act upon model insights.
- III. *Adaptability*: Solutions will need to accommodate an extensive range of models, designed for many use cases, and often with little insight into the training of those models. Every scenario cannot be planned for or built into the model. Tests are representative of a range of feasible scenarios to demonstrate how an individual is resilient and tolerant considering emerging unknown threats.
 - a. *Real-world Acuity*: Models must be assessed according to the trade-offs of mathematical performance and real-world use. Other metrics may need to be created.
- IV. CalypsoAI recommends that the NAIRR use TEVV to assess robustness, security, reliability, and bias of AI/ML algorithms during development and deployment. Examples of ways to assess AI/ML performance are as follows:
 - a. Model evaluation

- i. Demonstration of best practices used in the machine learning lifecycle
- ii. Recommendation on thresholds for acceptable/tolerable performance
- iii. Accuracy, Recall, Precision, Specificity
- iv. ROC Curve
- v. F1 Score
- b. Verification and Validation measures
 - i. Model simplicity: model based – Measure to determine how simple or complex a model is
 - ii. Model simplicity: feature based – Evaluation of the model performance relative to the percent of input features
 - iii. Noise injection – Measure to determine how much of the data used in the model is corrupt or not understood by the system (i.e., outliers, randomness)
 - iv. Bias Detection / Fairness Evaluation– Evaluation of the correlation between biased/unethical features, such as race and gender, and the other features used to train the model
 - v. Parameter modification – List of what model parameters were modified independent of the training data to control the learning process and justification for doing so
- c. Way to validate AI/ML performance test results
 - i. Proof of data spreadsheets used
 - ii. Demonstration of the model using new data to showcase performance matches the results of the previous testing

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private sector activities, resources, and services?

Government, academia, and the private sector have all recognized the need to harness the volume, velocity, variety, veracity, and value of data. As a result, they have invested significantly into various activities, resources, and services for artificial intelligence to compute data. For example, in 2018, the Defense Advanced Research Projects Agency (DARPA) announced a \$2 billion multi-year investment in AI projects that has included R&D efforts into High Performance AI.² Over the summer, the National Science Foundation helped establish strong building blocks with universities through creating 11 new National AI Research Institutes, which includes the NSF AI Institute for Intelligent Cyberinfrastructure with Computational Learning in the Environment at the Ohio State University.³ Moreover, Georgetown University stood up the Center for Security and Emerging Technology (CSET), which is addressing a host of relevant issues related to

² [AI Next Campaign \(darpa.mil\)](https://www.darpa.mil/news-events/2018-07-26)

³ https://www.nsf.gov/news/news_summ.jsp?cntn_id=303176

AI, such as data and computational power and cybersecurity.⁴ Finally, in the private sector, there are countless specialized startups that can leverage their expertise to solve different aspects of the computing and data infrastructure problem.

While assessing data quality is important, it is equally vital to take active steps in tandem with these efforts to secure the data we have. Given the skill gap that we see across emerging technology sectors, risk mitigation features must be as easy-to-use and accessible as possible. Recognizing this, CalypsoAI provides users of various skill levels the ability to assess the quality of their models and make adjustments based on the factors that are most important to them. This can include adjusting for datatypes, environments, and natural or adversarial corruptions to determine model performance in real-world deployment scenarios. In effect, CalypsoAI can help the NAIRR protect privacy and, when applied to cybersecurity systems, democratize cybersecurity infrastructure.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public-private partnerships should be central to the NAIRR because the private sector brings innovative perspectives and technologies that would enable the NAIRR's goal of creating a holistic computing and data infrastructure. Some exemplars of productive partnership include leveraging the knowledge of Federally Funded Research and Development Centers (FFRDCs) and other think tanks with topical research expertise; drawing upon Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) opportunities that flow through innovation hubs; and implementing rotational programs to embed private sector knowledge into government.

Additionally, public-private partnerships should be used for advancing AI educational opportunities, particularly if it can work with universities to create a course or module geared towards improving the data and computing ecosystem. In the course, the NAIRR can work with private sector companies to expose students to emerging and innovative capabilities. For example, since CalypsoAI's AI Model Risk Management platform, VESPR™, can be used by non-technical individuals, the NAIRR can leverage it as a teaching tool to broaden access to AI research and education opportunities.

Finally, the Task Force could again refer to the FLC as a model for public-private partnerships. For example, FLC offers a comprehensive database highlighting funding resources for federal agencies, academia, and industry. FLC also implemented the SBIR/STTR programs mentioned earlier, which have been widely successful in funding emerging technologies and R&D from small businesses and universities. If the NAIRR does not directly leverage SBIR/STTR programs, it could follow suit with its data sharing database and create similar programs to encourage public-private flow of ideas, engagement, and technologies.

⁴ [Publications - Center for Security and Emerging Technology \(georgetown.edu\)](#)

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

There are a few limitations that could potentially affect the NAIRR's ability to democratize access to AI R&D.

One limitation is the Task Force's ability to identify the specific data types and specific tools needed for the NAIRR toolbox. For guidance, the Task Force should look to industry for the most common and sought-after data types and tools.

Second, balancing funding and demand within the NAIRR is challenging because the cost of purchasing the correct data, computing power and processing, and software can be expensive. To address this issue, there should be task force liaisons responsible for engaging industry, government, and academia for specific topic areas such as Test and Evaluation (T&E), Verification and Validation (V&V), and data quality to assess demand. Additionally, to balance funding and demand insight can be drawn from partnerships with industry, annual budget analysis and adjustments, and utilization of existing and complementary resources to maximize funding.

One of the biggest limitations of implementation is the lack of common understanding and universal language between AI policy makers, users, stakeholders, and decision-makers. Without a common understanding between actors or universal language, AI-related risk cannot be properly managed. AI jargon is often not understood by key decision-makers, who assess the real impact of the risks associated with AI and are responsible for determining the level of acceptable risk. Furthermore, AI teams and vendors typically work on a discreet part of an overall operational problem. They target AI-specific metrics, such as high performance or accuracy, without necessarily understanding the complete context in which models will run. Therefore, the NAIRR should work in conjunction with other government agencies, industry, and academia to define universal language as it relates to AI.

Finally, AI/ML initiatives continue to be siloed, with isolated actors attempting to build and deploy models with limited access to standard industry tooling. This often results in inefficiencies and redundancy in AI efforts, as well as a lack of communication across organizations on best practices. The NAIRR can be one way of addressing these siloes through a collaborative effort, bridging the gap between government, industry, and academia. However, it will also require a great investment of resources from all stakeholders. As resource investment increases, greater regulation and oversight is needed. Sound oversight includes established evaluations, measures, and considerations for areas of interest such as accuracy, explainability and interpretability, reliability, privacy, robustness, safety, security, and mitigation of unintended and/or harmful bias. These all contribute to ensuring accountability and AI model robustness which, although more than a checklist, creates parameters for established good practices.

Conclusion:

CalypsoAI firmly supports NSF and OSTP's effort in establishing the National Artificial Intelligence Research Resource and appreciates the opportunity to provide our thoughts and feedback on the path forward. We welcome any opportunity to work with the Task Force, industry, and broader government agencies to assist in developing an accessible, inclusive, responsible, trustworthy, and secure NAIRR for the benefit of all sectors.

For further questions or for more information please do not hesitate to reach out to Hannah Mezei at [REDACTED]

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Ben Freed and Howie Choset

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

The 3Ds of AI: Data, Developer and Democratization

Ben Freed and Howie Choset

School of Computer Science, Carnegie Mellon

Abstract

As our AI tools become more advanced, they are increasingly created and controlled by a select few organizations. By limiting the breadth of institutions, groups, and people who can create, use, and inspect AI tools, the AI oligarchy has negative impacts for individuals, society, as well as the progress of the field. We firmly believe that AI has the potential to bring a tremendous amount of good to the world, but only if developed and used responsibly, which is a conversation in which everyone should have a voice. We identify three key ingredients to advance AI, which we call *the three D's of AI*: data availability, developer accessibility, and democratization of AI tools. It is our view that to unlock the true potential of deep learning, and retain American competitiveness, we must conquer the three Ds of deep learning: data, developer, and democratization.

The Three-D's	Data availability	Developer	Democratization
Benefit	Benefits everyone Removes biases	More talent Uncork talent	Society engagement Equitable ownership
Problem	Cost to create Privacy Biases / silos	Cost to obtain Cost to maintain Lack of tools	Unchecked firms Monopoly / barriers No easy-to-use tools
Solution	Surrogate data Public funds reqs. Data efficient apps	Tool develop Shared resources K12 education Retain international	Easy-to-use tools User-owned data K12 competition

Introduction

In recent years, AI has seen a boom of development, spurred by deep learning. Deep learning has revolutionized the way in which artificial intelligence is applied to domains such as manufacturing, finance, medicine, energy, agriculture, security, retail, just to name a few. Deep learning can be viewed as a type of data science that can model and predict future outcomes from (an enormous amount of) data that is provided to it, during a training process, say of a multi-layer neural network. Deep learning technologies are typically classified as *data driven*, because they primarily focus on extracting patterns from data, rather than relying on the knowledge of AI engineers or domain experts. This

shift in perspective from the *good old-fashioned AI* (GOFAI) techniques of past decades has the benefit that it removes human bias from the system, allowing deep learning algorithms to discover their own data representations and decision-making procedures, yielding higher levels of performance compared to hand-engineered approaches.

The salient feature that delivers deep learning's greatest strength- its ability to process an enormous amount of data - is also a drawback: it requires an overwhelmingly large amount of data to be effective. Such data may not be available to the "common" developer. In fact, lack of access to data is just one barrier of entry to enjoy the benefits of deep learning: an extensive and often time-consuming and expensive education is another requirement and therefore limits deep learning to highly educated and specialized PhDs with years of important education and training. These PhDs are great, but only represent the tip of the iceberg of potential developers that can contribute to and benefit from deep learning - not everyone can get into Carnegie Mellon. Finally, the computational resources to develop and use deep learning tend to be limited to the Google's, Facebook's, and perhaps some Universities of the world and yet many can contribute, if resources or low-overhead deep learning approaches were available.

An increasing portion of AI breakthroughs are being made using resources far outside the budget of the typical academic lab. Freelance and small companies also offer us opportunities that large companies and universities cannot, such as niche applications of AI to problems that might not be appealing to large companies; we do not want to lose them. Finally, improving access to data and AI tools also has the potential to reduce harmful bias in our AI technologies. If datasets are free and open, they can be inspected and are open to criticism by experts in fairness and ethics in AI. America may be at the lead of AI research, development and use, but frankly we are still doing it with one arm tied behind our back.

D1: Data availability

Deep learning-based approaches require a large amount of data to be effective. As can be expected, data availability plays a crucial role in the performance of our ML-based technology. High-quality datasets are necessary for the advances made in research settings to percolate into applied technologies, because availability of quality data plays a large role in determining the efficacy of a learned model in the real world. While it is a strength to process and "learn from" a large amount of data, often high quantities of data is required for development and training of AI approaches. Typical datasets used to train deep neural networks used for supervised vision models contain hundreds of thousands to millions of labeled datapoints (e.g., ImageNet, CIFAR10, CIFAR100). For example, state of the art natural language processing (NLP) models, such as GPT3,

have been trained on hundreds of billions of words. State of the art reinforcement learning systems trained for two-player game-play, such as AlphaZero and AlphaStar, were trained on games. For typical academic labs, these large data requirements limit the applicability of deep learning to situations in which large datasets are freely available.

Challenges limiting data availability include 1) high cost or high required investment, 2) privacy concerns, 3) data is siloed and 4) difficulty in obtaining high-quality labels.

D1P1: High cost / investment to acquire data: Gathering large datasets for custom applications is often a high-cost endeavor. For example, one recent publication that used reinforcement learning to learn robotic grasping required 14 expensive robotic arms and over 800,000 grasp attempts to achieve 80-90% grasp success rates using a 2-finger gripper [CITE Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection]. For many academic labs, this would be a prohibitively expensive undertaking. Finally, in order to get the most out of our taxpayer dollars, we must share and document our publicly supported datasets.

D1P1.1: Data efficiency: To ameliorate the high costs associated with large dataset collection, we suggest that more funding be allocated to improving data efficiency in ML, through techniques such as transfer learning, semi-supervised learning, domain adaptation, and data augmentation. These approaches improve data efficiency on some target task by enabling data from another task, or unlabeled data, to contribute to the learning process.

D1P1.2 Leverage public funds: We can also leverage our existing investments by requiring that datasets generated by public funds should be available to the public: it was paid for by the public. Also, we believe that such a practice will inevitably promote the scientific process of validating results. In fact, we already see this practice taking place, as it is in the best interest of the scientist to promote their work and supporting data facilitates such promotion.

D1P1.3: Synthetic Data: Finally, we advocate for increased funding for research on synthetic data creation. Fake data is free but making it meaningful is hard; however, the ability to generate realistic synthetic data could at least allow AI researchers and practitioners to validate approaches and identify weaknesses before making costly investments in dataset collection. An additional potential benefit of synthetic data generation is that it avoids the privacy concerns typically associated with sensitive (e.g. medical) data.

D1P2: Privacy. Privacy permeates all issues that involve datasets. Obviously, privacy at the extreme inhibits the proliferation of datasets in order to protect the owners and subjects of the data. Failure to recognize this importance could be catastrophic. For one thing, we can compromise personal, organizational, and national security. Next, we could potentially lose the trust of those who contribute to the dataset. By no means do we, the authors of this document, claim to be privacy experts, nor understand the bounds of the implications of privacy. Therefore, we suggest that experts in privacy be included in the ideation of data availability and defer to them for specific suggestions. With such experts, we advise that approaches be developed to either develop surrogate data sets and other methods be created to strip private information from datasets and yet retain their salient properties.

D1P3: Data Silos. Improving data availability has the potential to both improve AI both as a research field and a technology. Data is the raw ore from which useful models are smelt, and machine learning is a fundamentally empirical science; hypotheses must be tested on *real-world data* that are truly reflective of the situations in which they are intended to be used. For this reason, in many fields such as computer vision and natural language processing, large, representative, and high-quality datasets are absolutely crucial for fundamental advancement. We strongly advocate an increase in funding for transfer learning and imitation learning, but require disparate problem domains for which this research would be funded.

Additionally, we suggest that measures be taken to encourage inter-agency sharing of data, when possible. It stands to reason that different agencies may have some common denominators in the datasets they collect. It would be meaningful to understand the commonalities, to see what shared problems they're all solving, as well as the differences to see how we can round off each others' limited datasets. Moreover, agency x can stress-test its approaches using agency y's datasets. The problem is that, from the authors' distant perspective, agencies often have a hard time cooperating and sharing at deep levels. We suggest that the White House look at examples of where inter-agency cooperation has been successful, and one such example is the National Robotics Initiative, based out of the NSF.

D1P4: Labels

Labels typically refer to some form of identification or annotation placed on data by people. Obtaining high-quality labels can in many cases be the most expensive aspect of data collection. Often, gathering unlabeled data is cheap because it requires little human oversight (e.g., downloading text from wikipedia or images from Google images). In some applications, such as labelling of medical data or data from particle accelerator experiments, data must be labeled by domain experts, who's time is very valuable.

D1P4.1: To ameliorate the difficulties associated with labeling large datasets, we advocate for increased funding in dataset generation. Generation of high-quality quality datasets with high-quality labels is not a flashy job, but it often spurs advances in the field (e.g., the ImageNet dataset, which was an expensive undertaking, but since its inception has served as an invaluable tool for the computer vision community). Of course, we must acknowledge that incorrectly labeled datasets can have a detrimental effect, but

D1P4.2: We additionally suggest increased funding for machine learning approaches that make more efficient use of human experts, for example *active learning*. Active learning is a form of machine learning that allows the ML system to query an expert or other knowledge source (e.g., a person, or a simulator) for labels during the learning process. Typically, active learning algorithms are designed so as to query the expert for the highly useful information, thereby reducing the number of labels that must be provided by the expert.

D2: Developer Access

Developer access relates to the resources and capabilities that people who develop AI technologies must possess in order to develop, and in many cases advance the state of the art, in AI and deep learning. One resource, as described above, is data. However, other resources are needed: computers, software tools, developer communities, etc. Just like data, a tremendous problem faced by deep learning developers is the quantity of computational resources and other developer access tools required. Most academic labs cannot compete with the massive GPU (and now TPU) clusters used by the likes of OpenAI and Google. As a result, high-powered private industrial companies such as Facebook and Google, would be the only ones who could enjoy the benefits of developer tools to advance the state of the art. This means that the most powerful AI algorithms are controlled by a few large companies. We should seek a policy of supporting research and education in empowering people outside these centers of machine learning excellence to create novel AI tools.

D2P1: Computing resources. The primary obstacle to developer accessibility is cost. The hardware cost for a single AlphaGo Zero system in 2017, including the four TPUs, has been quoted as around \$25 million (according to wikipedia). Certainly, a tier 1 University lab, let alone a small company or citizen-scientist, cannot afford such computational resources. The trend towards ever larger models that yield better performance on popular benchmarks while requiring more computational power to train

makes it increasingly difficult for labs with modest resources to compete with state-of-the-art (SOTA) performance on ML benchmarks.

D2P1.1: Shared Resources. Therefore, we suggest, just as the physicists can band together to raise funds for a common platform, such as a telescope, so should the academic AI researchers form a similar consortium for a shared resource. This could follow the already existing model of the Super Computer Centers, but some careful consideration must be given to the special needs of AI researchers and perhaps the more broad user community of such a resource.

D2P1.2: Efficient tool development. To better enable labs with smaller budgets and modest compute resources to compete with well-funded companies, we recommend that the NSF fund research in *computationally efficient* machine learning approaches, and *low-cost computing hardware (perhaps including robotics)*. Investing more in computationally efficient ML approaches would have several benefits: firstly, it would provide more avenues of possible funding for labs that are capable of contributing, but cannot match the SOTA performance on benchmarks simply due to computational limitations. Secondly, the development of computationally efficient ML approaches would allow academic labs with modest budgets *to be competitive* with SOTA performance. Finally, efficient ML algorithms have the potential to lower the carbon footprint of ML research.

D2P2: Another challenge limiting developer accessibility is K-12 education. Opportunities for K-12 students to engage with computer science are not evenly distributed, putting segments of the population at a disadvantage when entering college. We therefore advocate for the expansion of computer science education in K-12. Computer science is unique in that compute resources, and even IT support for students can be easily shared by multiple schools.

The barrier to entry for becoming an AI developer for the community at large is unnecessarily high. Even state-of-the-art advances in machine learning can be broken down into a few basic steps: define the model, train the model, validate the model. While programming libraries exist (e.g., Keras) that massively streamline the machine learning development process, even just installing and using these tools requires a high degree of programming expertise and understanding of computer infrastructure, for example, proficiency in python and linux.

To lower the barrier to entry for potential AI developers, both in K-12 and in the community at large, we advocate for the creation and development of web-based tools, accessible to anyone with an internet connection, that allow machine learning workflows

such as data set handling, model definition, and training, to be represented through a simple and easy-to-use graphical interface. Any program created in this interface could then be converted to (e.g. python) code for the purposes of further development or sharing with the AI community.

D3: Democratization of AI Tools and Data.

Beyond data and developers, AI tools are often out of reach of most people who want to use AI tools to solve problems for their own businesses, or just personal research and education. For the United States to reach its full potential in using advanced computing to compete and collaborate with our peers in Europe and Asia, we must get the AI solutions into the hands of everyone. We believe that in doing so, everyone has an opportunity to voice how AI tools are used - in other words, we must democratize the use of AI tools.

As stated already, having solutions in the hands of a few large companies will limit our ability to solve complex problems. We are quite fortunate to have Tenosor Flow and Pytorch, owned by Google and Facebook, but as AI tools increasingly shape our lives, it is increasingly important that the power of tech giants does not go unchecked. Having the citizen-AI-scientist using AI tools to solve similar problems may actually serve as a check and balance to large companies, whose initial goals were to generate profit, who may misuse or abuse their capabilities.

One major challenge toward democratizing AI tools is the fact that large tech companies (such as Apple, Google, and Facebook) control much of the data generating pipelines, because much of the data these companies run on is generated by users using their products. While these tech giants offer a tremendous benefit to our economy and society, we cannot allow them to monopolize the AI market in perpetuity. We are inhibiting our growth if we sustain long-term difficulties for small companies or non-profit open-source ventures to break into the market. This challenge overlaps heavily with the issues discussed in data availability; however, here we are mainly focused on the assumption that companies own the data generated by their platforms, and how this limits democratization of AI tools.

To encourage competition in the AI market, as well as decouple data from data-generation platforms, we advocate for measures to be taken that allow users of AI technology to *own their own data*. Users could then choose to sell their data at free market price on a data market, thus lowering the barrier to entry of smaller companies and research groups. Such a data ownership model would also give users the ability to

vote with their data: if users do not like the way a particular company uses their data, they can choose to withhold their data from that company. Changing the data ownership model gives users of AI technology a seat at the table, instead of simply allowing big tech companies to be the sole arbitrator of how to use data and AI tools in whatever way makes them the most profit. Finally, it is our belief that having a free data market will have the positive side-effect of encouraging citizens to use and develop AI tools.

Federal Register Notice 86 FR 46278, <https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Carnegie Mellon University

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Carnegie Mellon University

**Submitted by Carnegie Mellon University
In Response to the Request for Information on the
Implementation Plan for a National Artificial Intelligence Research Resource**

Contributing authors include Howie Choset, George Darakos, Motahhare Eslami, Matthew Gaston, Michael McQuade, Eric Nyberg, Norman Sadeh, Shane Shaneman, Russell Schwartz, and David Thompson

Introduction –Seizing this Critical Moment

Carnegie Mellon greatly appreciates the opportunity to share perspectives on the development of a National Artificial Intelligence Research Resource (NAIRR). The work of the Task Force will help set the foundation for the United States’ continued leadership in AI research, development and deployment. The high cost of computing for a single large-scale deep learning system can reach into the tens of millions of dollars. This resource will also play a critical role in ensuring that our nation realizes the potential for advances in artificial intelligence to address major challenges and achieve vital national aspirations, and it will inspire and enable the contributions of the broadest, most diverse set of American talent.

Recognizing the significance of this moment, this response seeks to underscore the following underlying tenets:

- The development of a NAIRR should build upon the lessons of successful past federal investments in national computing and science research infrastructure, particularly models that have demonstrated an ability to contribute to goals such as the regional diversification of innovation and which have fostered both strong public/private collaboration and community-building.
- As articulated in the RFI, the NAIRR strategy must consider the multiple national missions that AI is being called upon to help address. They include ensuring sustained U.S. global leadership in critical technologies, meeting the climate challenge, enhancing the resilience of public health and national supply chains, and advancing a more equitable society by restoring the power of innovation as a force for social mobility and opportunity across all communities.
- This effort is vital to ensure access to computing resources, as advancements have increased costs and created barriers that impede research. Computing requirements associated with technologies such as deep learning are massive. Large tech companies have access to computing power that far exceeds resources available to typical academic research groups. An

infrastructure that makes similar resources available to the academic AI community will significantly enhance our national research capability in AI.

- Yet, as the RFI also recognizes, this challenge requires investments in, and strategies supporting, the development of human capital—particularly impactful efforts to foster a more diverse science community.
- Finally, the NAIRR must also be designed to draw upon and catalyze advances at the frontier of AI research, advances that that will shape future computing infrastructure demands and innovations.

The creation of a National AI Research Resource comes at a critical inflection point in the evolution of strategies to shape the future of America’s scientific research infrastructure. It also coincides with bipartisan support for historic investments in science, including basic and mission focused research and the development of testbeds and other infrastructure designed to accelerate the pace of innovation. Perhaps most profoundly, it coincides with advances in AI that hold the potential to transform the research enterprise itself through automated and AI-enabled scientific discovery.

Thoughts and Recommendations

The following recommendations reflect the view that the development of an effective NAIRR involves building cohesive and appropriate strategies at three key layers. In addition to its primary focus on computing infrastructure the NAIRR will need services and capabilities at data and human capital layers.

Question One: Building the Roadmap for NAIRR

Capitalize on the National Science Foundation’s Legacy of Leadership and Emerging Capabilities.

We embrace the focus of the National AI Innovation Initiative and recommend capitalizing on the strengths of the National Science Foundation to serve as the lead agency for the NAIRR. NSF has a history of aligning investments in infrastructure with initiatives to support career development and broad engagement. NSF has also advanced model public/private collaboration in supporting both research and advanced scientific infrastructure. This lead role will also strategically align with initiatives to advance mission related research at NSF and more focused initiatives in diversity, equity and inclusion and a more comprehensive focus on STEM education.

Leadership of the NAIRR will also demand strong inter-agency coordination. NSF has also advanced successful models of multi-agency engagement, such as the National Robotics Initiative (NRI), upon which it can draw to help inform strategies for the NAIRR. Building upon the NRI model, NSF could structure a multi-agency coordination team to help guide the NAIRR that will facilitate shaping a shared national vision and strategy while also ensuring that the distinct missions of each agency are advanced.

Draw upon Successful Models of Building Regional Hubs

Important lessons can also be drawn from an earlier generation of investments in advanced computing infrastructure. While the technical challenges and critical mission objectives are distinctly different and will be supported with a cloud-based infrastructure, the creation of NSF’s original network for

supercomputing centers offers important insights of value to the NAIRR effort. There, the creation of a network of competitively awarded regional centers fostered the emergence of a diverse set of technical and research management capabilities that included the rise of complementary areas of specialization in different centers. This network of centers also accelerated the development of university partnerships with industry while encouraging advances innovation driven by competition. Additionally, these centers can serve as a focal point for outreach and diversity, offering grants and educational experiences to those who otherwise would not have easy access.

The network of supercomputing centers approach also leveraged state and local investment, enhanced the linkage between advanced computing capabilities and primary and secondary STEM education. This network model enabled NSF to advance a national mission and vision for these computing resources while capitalizing on local governance and operational innovations—including multi-institutional collaborations and industry partnerships. This network model will also lend itself to fostering multi-agency collaborations that align with the unique specializations of different centers to enhance innovation.

Finally, the regional nature of the centers contributed to expanding the geography of innovation. In the case of Pittsburgh for example, the award of one of the original supercomputing centers (the Pittsburgh Supercomputing Center) to an entity created by Carnegie Mellon and the University of Pittsburgh was an early catalyst to the growth of university based economic development in the region.

A NAIRR initiative will involve distinct technical challenges and investments and will necessitate the development of a broader and more complex set of support and operational services, but the model of creating a network of competitively awarded hubs will provide a valuable starting point for the NAIRR.

Question Two: Recognizing the Need for a Diverse Set of Services.

Provide Support Across the Project Continuum and Build a Broad Base of Services

The focus of the NAIRR is to advance access to resources needed for large scale computing research initiatives. In structuring services, the program should recognize that these research projects evolve along a continuum. Services should both support early stage and exploratory research activities and be capable of supporting long-term developmental engagements. While these early-stage projects do not demand the same level of support, they are still often impacted by difficulty accessing resources.

Build Services that Focus on Critical Challenges – Privacy, Ethics, Talent and Workforce Development and Democratization of Data

In addition to broadening access to computing resources and expertise in the management of deep learning applications, the NAIRR must provide a mix of services to support its mission of broadening access to AI research capabilities. These services should include advanced resources for the identification and proper acquisition of data, data curation, security, privacy and bias expertise. This focus should include expertise and provide technical assistance support to help institutions clarify and update IRB requirements and procedures.

The NAIRR should also build a network of collaborating institutions to assist in advancing education and training initiatives. Models for such collaborations that blend professional development, certificate and shared curricula across institutions have been developed in areas ranging from AI to cybersecurity and

defense engineering. A focus on building strong and broad collaborations in education and training can also enhance the mission of the NAIRR to democratize access to AI resources by enhancing capacity development.

A workforce strategy should also include a focus on helping advance models for building communities or networks of federal data users. These networks can enhance data access and security applications. One model that the NAIRR could build upon is the Coleridge Initiative, a not-for-profit organization bringing together researchers and government agencies to improve access, security and privacy and enhance the development of research initiatives. This model could be expanded to enhance data sharing with the private sector. Activities in this service area will require close coordination with the emerging data democratization initiatives of the White House Office of Science and Technology Policy and the work of United States Digital Service.

Finally, the democratization of access to computing capability should also provide a catalyst to accelerating AI innovation and entrepreneurship. Building collaborations with university tech transfer ecosystems and the emerging investments in regional innovation capabilities could facilitate the ability for the NAIRR to be resource for supporting both entrepreneurship initiatives and regional testbed developments.

Contribute to Advancing the Frontiers of AI

As the pace of innovation in emerging technologies continues to accelerate, it is critical to ensure that NAIRR AI infrastructure investments not only support the current application of AI to scientific problems and domains, but also accelerates the development of Distributed AI capabilities to further democratize access to AI data and resources across the United States. Carnegie Mellon has launched a strategic research initiative called *AI Fusion* that is focused on accelerating advances in Distributed AI to overcome the traditional AI paradigm which requires massive data sets to be aggregated and engineered centrally to have access to the most comprehensive AI algorithms and high-performance computing resources that enable AI and machine learning processing. Instead, AI Fusion interconnects multiple, disparate systems and AI resources as part of an immersive *AI Fabric* which then enables the fusion of data sets, compute resources, and AI processing distributed across the country. AI Fusion augments traditional AI capabilities by enabling a more robust and scalable nationwide AI infrastructure that is highly responsive and adaptive, providing much more capability at the 'point of need' while still ensuring synchronization and properly structured information sharing. Prioritized investment in AI Fusion and the accelerated development of Distributed AI are pivotal frontier investments in our future AI infrastructure.

This focus on distributed AI should be complemented by coordinated research initiatives that seek to directly address the critical factors impacting the growing cost, complexity and energy impact of deep learning. These targeted areas include a focus on Synthetic Data and initiatives to advance the generation of high-quality labeling.

The NAIRR should also focus on building the intellectual and technical infrastructure to advance multiple solutions for cloud interactions. An example of research in this area is the Army's Project COEUS. The project is building an AI ecosystem--a virtual location for users to gather for the optimization of sharing algorithms, uniformly stored and organized for accessibility. This ecosystem will provide tools to enable users to more rapidly integrate AI.

Contribute to Advancing the discipline of AI Engineering

AI Engineering is an emergent discipline focused on developing tools, systems, and processes to enable the application of artificial intelligence in real-world contexts. AI Engineering is a field of research and process development that combines the principles of systems engineering, software engineering, computer science, and human-centered design to enable researchers and practitioners to methodically develop AI systems that are robust, secure, scalable, and human-centered. To fully reap the benefits of the National AI Initiative and to bring the power of AI to our most critical national needs and our most exciting opportunities, development of the methodologies of AI engineering must be a core element of the NAIRR. Carnegie Mellon's Software Engineering Institute is focused on building models, tools and practices to advance this discipline. The NAIRR should build formal relationships with both academic institutions and Federally Funded Research and Development Centers to integrate AI engineering resources into its core capabilities.

Finally, the NAIRR should be focal point for accelerating the development and deployment of advances to improve the energy efficiency of computing. By helping foster distinct models of industry and university collaborations a network of regional resources centers may also help advance these innovations. This focus will necessitate collaboration across multiple agencies and engage academic talent across the nation.

The points outlined above highlight the extent to which the NAIRR must be able to galvanize and engage collaboration across the broader AI community. Recent AI initiatives have created mechanisms to help facilitate community engagement to advance strategic planning and initiatives. One approach would be to consider establishing an Academic Innovation Council. This model, which has been developed for DoD AI initiatives, has proven effective at drawing talent from across the nation to develop roadmaps to help inform future research initiatives.

Question Three: How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Build a Comprehensive Capacity to Integrate a Focus on Ethics and Privacy

In focusing on issues of security, privacy and bias, the NAIRR will have the opportunity to coordinate with and leverage NSF initiatives created through the enactment of the National AI Innovation Initiative and, potentially the final legislation emerging from the US Innovation and Competition and the NSF For the Future acts. In particular, the provisions of these measures focus on ethics training and the development of ethics impact statements as part of AI research proposal development can be integrated into initiatives of the NAIRR.

The NAIRR can contribute to these existing measures by creating a nexus for community engagement on these critical issues. The NAIRR could incorporate grand challenges and other programs that focus on innovations in tool development, curricula and community building into its service portfolio to contribute to and elevate engagement with broader federal policy efforts. Another approach that can reinforce principles of ethical and responsible research and development of AI is incorporating AI ethics into the curriculum of academic institutions. This could include catalyzing a national effort encouraging K-12--more schools to add AI & ethics into their programs. This can train knowledgeable future AI

researchers & developers who are equipped with the right tools to tackle potential ethical challenges of AI in their work. Finally, the NAIRR could work with the academic research community to develop model approaches for engaging broader community stakeholders in helping assess the potential impacts of AI research.

There is an urgent need to clarify and update the interpretation of IRB requirements in light of the collection and mining of human subject data and to ensure adequate awareness and training of AI researchers and IRB offices at universities. The NAIRR could take a leading role in coordinating this process, producing and possibly delivering relevant and effective training material, and being a central node for sharing best practices. The same applies to other ethical considerations, including issues of racial and gender equity, bias, civil rights, transparency and accountability - as well as issues of safety and security.

Another critical dimension of effective ethics and responsible research strategies is to advance capabilities for anonymization. Privacy and ethical concerns currently make providing high quality and representative datasets quite challenging. While different sectors, including private, academic, or public, rely on such datasets to train their AI models, they cannot share many of those datasets publicly as most of the anonymization procedures cannot completely remove the connection between the data and users. This requires thorough investigation of each dataset, reflecting on potential privacy issues, and coming up with mitigation plans. The NAIRR should be a focal point for advancing strategies to address these concerns

Question Four: What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

With the accelerating pace of innovation in AI, it is critical for the NAIRR to not only provide individual researchers access to critical research infrastructure and resources, but also to facilitate collaboration of the best and brightest researchers across the country. As part of Carnegie Mellon's long-standing and strategic relationship with the Army Research Laboratory (ARL), we worked closely with ARL to develop a collaborative portal that will dramatically enhance collaboration with the best and brightest researchers across the United States while also giving them secure access to tools and resources, and a virtualized computing environment to accelerate AI research with government trained models or data sets. ARL worked closely with Carnegie Mellon to develop an on-boarding process that is academic-friendly and does not restrict access to the collaborative portal based upon citizenship status or require researchers to qualify for a DoD Common Access Card (CAC).

Once on-boarded into the A2I2 Collaborative Cloud, researchers are able to send and receive CUI emails, take part in virtual Microsoft Teams meetings, and collaborate with other researchers, industry partners, and government personnel. The A2I2 Collaborative Cloud allows researchers across the county to have a virtual CUI desktop and research space at a DISA Impact Level 4 (IL4) level using their existing University PC / laptop and network. Enabling US and foreign researchers the ability to have ready access to government data and CUI information, drastically enhances their ability to develop a targeted approach to research with a much better understanding of the problem space, which in turn drastically increases the impact of AI research. This collaborative approach has been leveraged extensively by the private sector as well to accelerate AI research.

Some additional potential resources that the NAIRR can leverage include the existing frameworks and infrastructures some private-sector companies such as Google (People and AI Research: PAIR) and Facebook (Facebook AI Research: FAIR) have built around AI and research. These groups usually have a well-documented structure, guidelines, and resources for researchers who want to do AI work.

The NAIRR can also leverage existing educational frameworks that have been developed to support and enable the future AI U.S. workforce and expose more people at a younger age to research tools. For example, CMU has developed educational infrastructure at the K-12 level through post graduate levels to educate students at scale. Some of the following programs could potentially be leveraged to incorporate curriculum that could expose students to NAIRR through all educational ranges:

K-12:

- Computer Science Academy (CSA) - a free, universally accessible, online, interactive high school computer science curriculum designed and managed by CMU undergraduate students. CSA offers teacher training, an online interactive textbook, and online technical support from undergraduate computer science students, available “24/7”.
- Computer Science STEM Network (CS2N) - The is a collaborative research project between Carnegie Mellon University, including the Robotics Academy, and the Defense Advanced Research Projects Agency (DARPA) designed to increase the number of students pursuing advanced Computer Science and Science, Technology, Engineering, and Mathematics (CS-STEM) degrees.

Community College

- SAIL - an online learning platform that provides college and university instructors with job-focused technology courses created at Carnegie Mellon University that are project-based, collaborative, and use real-time feedback.

Similarly, the NAIRR can also leverage university based models that foster collaboration among researchers from technical, policy, business and the humanities to engage on critical issues related to the development and impact of AI. For example, Carnegie Mellon’s Block Center for Technology and Society brings such an interdisciplinary focus to research on bias and ethics and the use of AI to address major societal issues. Block Center projects also help provide a bridge between the research community and government users of AI applications to foster best practices and build capacity for the ethical deployment of AI.

Question Six: Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The democratization of AI is central to the mission of the NAIRR. Thus, many of the comments above focus on elements of this democratization objective. These elements include the imperative of building services and capabilities at the infrastructure, data and human capital layer as well as proposals to create a vehicle for input from a broader stakeholder community and the imperative of building a focus on ethics and privacy that leverages emerging policies and programs at individual institutions and the federal research agencies. A focus on creating regional hubs as a central element of the NAIRR approach can also contribute to this democratization initiative by helping expand the geography of innovation.

Ultimately, a critical variable will be the need to focus on democratization objectives throughout the metrics that will be used to judge the value of the NAIRR operations and outcomes. A focus on metrics will help ensure that the NAIRR is a catalyst to fostering community wide engagement and commitment to this fundamental goal and that this commitment is reflected in all projects and collaborations.

Concluding Thoughts

Carnegie Mellon University is synonymous with the birthplace of artificial intelligence and continues as a critical hub defining the future of AI and, through it, the future of society. Our faculty and students have been at the forefront of the development of AI technologies that have changed how we live—from self-driving cars to personal assistants, robotic surgery, cognitive tutors, AI enabled traffic signals and the applications of machine learning to fight food scarcity. Carnegie Mellon is committed to continuing to advance this frontier of research and education. The development of a National AI Research Resource is a vital step to ensure that the U.S. leads the development of AI and that Americans in all communities benefit from the potential of these advances to expand economic opportunity and improve the quality of life.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Center for Data Innovation

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



1. What options should the Task Force consider for any of roadmap elements A through I, and why?

[A] Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success.

There are two broad goals the Task Force has already rightly identified a NAIRR can help achieve: More AI innovation and bolstered U.S. competitiveness in AI.³ However, it is important for the Task Force to recognize that even though many use these terms interchangeably, AI innovation and competitiveness mean different things and optimizing the NAIRR implementation roadmap to maximize for one can lead to different solutions than maximizing for the other.

First, consider how these terms differ. Competitiveness refers to the ability of an economy to compete effectively in global markets for traded goods and services in the absence of subsidies and government protections.⁴ By enabling an economy to export more in value added terms than it imports, competitiveness increases a nation's standard of living. In contrast, innovation refers to developing an improved product, production process, or organizational method. If this innovation occurs in traded sectors, a nation's economy will become more competitive. But innovation in non-traded sectors will have less impact on competitiveness because by definition their output is not sold outside local borders.

These distinctions matter for implementing, operating, and administering a NAIRR. To see why, consider two AI researchers, one of whom is pursuing research into AI models for construction and the other for manufacturing. Both research pursuits support AI innovation but only the latter would bolster U.S. competitiveness in AI because manufacturing is a traded sector from the perspective of the U.S. economy while construction is not. A NAIRR whose primary goal is to promote AI innovation should seek to support both types of research equally whereas a NAIRR whose primary goal is to bolster U.S. competitiveness should prioritize AI research for manufacturing over construction.

To be clear, boosting access to research tools can serve as means to both ends. But the Task Force should consider having well-articulated and distinct mechanisms to achieve each. One way to do this is by implementing separate support mechanisms for academic

³ The White House, "The Biden Administration Launches the National Artificial Intelligence Research Resource Task Force," news release, June 10, 2021, <https://www.whitehouse.gov/ostp/news-updates/2021/06/10/the-biden-administration-launches-the-national-artificial-intelligence-research-resource-task-force/>.

⁴ Robert D. Atkinson, "The Competitive Edge: A Policymaker's Guide to Developing A National Strategy," (ITIF, December 2017), <https://www2.itif.org/2017-competitive-edge.pdf>.



researchers and private sector researchers to access the NAIRR. For instance, one mechanism could provide support for—and only for—eligible academic and government researchers who are conducting AI research that promotes AI innovation in any field, with research proposals reviewed through a competitive process. Another mechanism could provide eligible firms with innovation vouchers they can use to “buy” AI compute time and expertise at certain supercomputing centers, with the size and type of the voucher determined by its relevance to a national competitiveness strategy (e.g., focused on solving specific challenges and facilitating commercialization breakthroughs).

The role of government in increasing access to AI resources for academic and private sector researchers are different. Academic researchers typically conduct crucial early-stage AI research that provides foundational, generic knowledge that everyone—including industry—can draw on for ideas and innovation. However, only well-resourced institutions provide access to expensive AI resources, such as powerful AI compute. The government’s role is to ensure as many qualified academic researchers as possible have access to AI resources in order to expand the pool of general AI knowledge for the benefit of everyone. Private sector researchers typically conduct later-stage R&D, which is important in bringing innovations to market. The private sector already has incentives to invest in AI resources. The role for government is to ensure the private sector’s incentives to invest in R&D for AI are sufficient to maximize overall economic welfare.

The Task Force has also rightly recognized that democratizing access to AI compute for academic researchers can help ensure all individuals have equal opportunity to succeed in becoming the next generation of AI researchers. The Task Force could introduce a third support mechanism that specifically supports the allocation of resources at Minority-Serving Institutions (MSIs) that include Historically Black Colleges and Universities, Hispanic-Serving Institutions (HSIs), and Tribal Colleges and Universities (TCUs) to help achieve this end.

[B] On a plan for ownership and administration of the National Artificial Intelligence Research Resource.

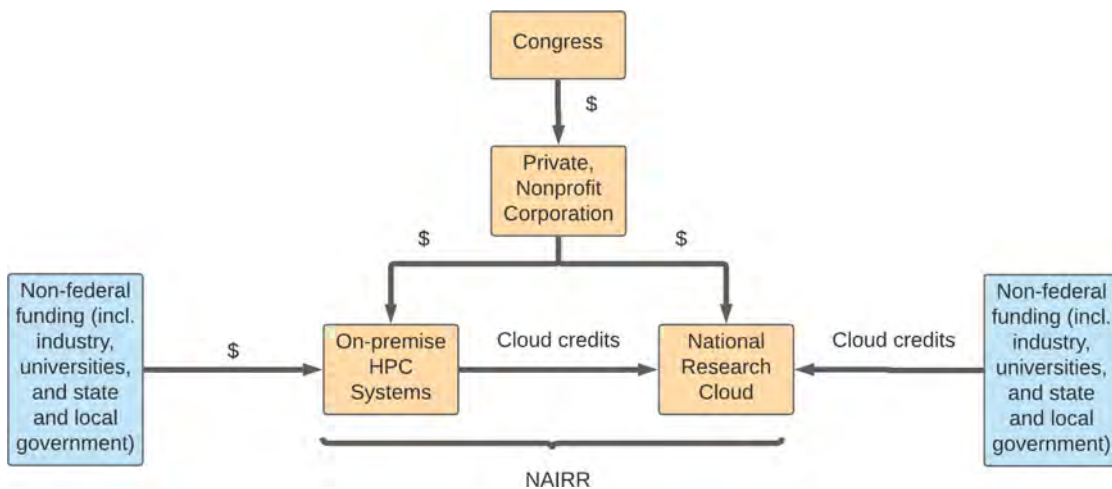
One option the Task Force should consider is for a private, non-profit corporation to allocate federal funds to the NAIRR. This entity would be created by Congress and act as a steward of the federal government’s investment in national AI research resources. To see how this might work, consider the Corporation for Public Broadcasting (CPB), a private, non-profit corporation established by the Public Broadcasting Act of 1967 to act as a steward of the federal government’s investment in public broadcasting. The mission of the CPB is to “ensure universal access to non-commercial, high-quality content and telecommunications



services.”⁵ It does not produce or distribute programs, nor does it own, control, or operate any broadcast stations. Instead, CPB allocates federal funding to local radio stations in each of the 50 states, which broadcast national content to their local communities and broadcast local programs they create themselves too.

We propose the Task Force consider a comparable model for the administration of the NAIRR. Through the appropriations process, Congress would enact federal payments to a private, non-profit corporation, just as it does for federal agencies that fund high-performance computing (HPC) at supercomputing centers and universities (see Figure 1). The corporation would not own, operate, or control any HPC systems itself but instead be charged with facilitating geographic diversity of AI compute, the development and expansion of HPC for AI, and providing funding to local HPC systems.

Figure 1: Proposed model for the NAIRR.



The activities of the corporation could be twofold: 1) to allocate federal funds to local, on-premise HPC systems at universities, colleges, and research institutes across the country; and 2) to provide funding for nationally accessible resources such as a National Research Cloud.

Regarding the former, local systems could be owned and maintained through public-private partnerships, which is discussed further in section 4. And to understand the latter, let us

⁵ “About CPB,” last accessed September 10, <https://www.cpb.org/aboutcpb>.



return to the example of public broadcasting. National programming producers like NPR, APM, and PRI are independent entities that are funded through a number of sources including corporate sponsorships, funds from CPB, and fees from locally owned and operated radio stations that pay to be their members and distribute their programming. A similar set up could work for nationally accessible HPC as part of the NAIRR. For instance, one nationally accessible resource could be a National Research Cloud (NRC), set up as publicly and privately funded non-profit with member institutions across the country. The members (both public and private) would make some level of AI compute available in the cloud and gain access to government datasets and other incentives from the NRC in return. Because it would be a public-private non-profit, the NRC could partner with private companies to obtain cloud services from existing vendors for AI researchers, which would be particularly valuable in the short-term as it gets established. In addition, local institutions would have a choice. They could choose not to participate in the NRC and exclusively provide local, on-premises AI compute, which will be important for some researchers who require on-premises resources for reasons such as data security, application performance, or teaching purposes.

Such a set-up would be adaptable, allowing for the incorporation of new resources and novel computing capabilities. One important and related question the Task Force raised in a recent workshop was which regions should it prioritize for on-premises systems?

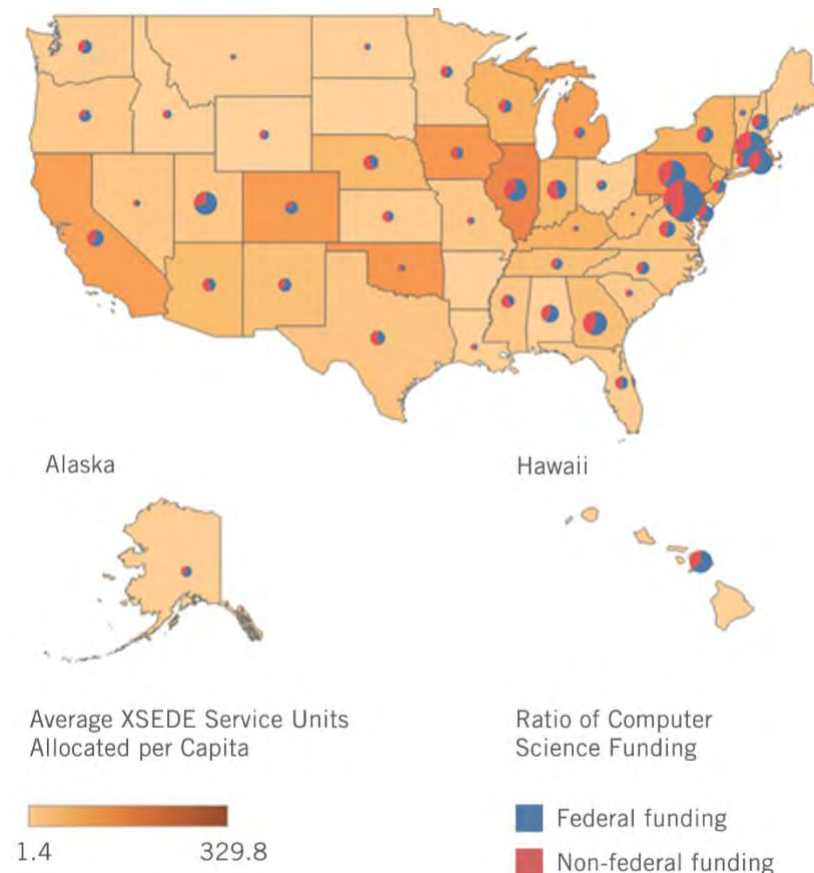
Which regions should the Task Force target for providing local systems?

The NAIRR should prioritize providing local resources in regions wherein the gap between AI compute demand and supply is greatest.

Some communities, institutions, and regions already have high access to HPC availability while others are conducting high levels of AI research but have little access to powerful systems. In our 2020 report *How the United States Can Increase Access to Supercomputing*, we provided an estimate of access to HPC resources per capita across the United States. We used data on compute time researchers requested in 2017, 2018, and 2019 from NSF's Extreme Science and Engineering Discovery Environment (XSEDE), a platform that coordinates the national sharing of supercomputing, as well as data on the researcher's organization and the state in which the organization is located (see Figure 2).⁶

⁶ Hodan Omaar, "How the United States Can Increase Access to Supercomputing," (Center for Data Innovation, December 2020), <https://www2.datainnovation.org/2020-how-the-united-states-can-increase-access-to-supercomputing.pdf>.

Figure 2: Proportion of XSEDE service units allocated per capita and size of research funding from high-research institutions in computer science per capita in each state.



The key insight from this figure is more access to powerful HPC resources is found in states like Massachusetts, Pennsylvania, and Illinois that have leading academic institutions, which can either stand up their own HPC centers or partner with other leading research institutions in their state to create multi-institutional centers.⁷ Federal investments in more HPC resources in regions where HPC availability is already high will not be the most effective way to close the gap between HPC demand and supply because institutions either already have baseline AI compute and are using it for research, or they don't have research funding which means access to HPC is not the problem, research funding is.

⁷ Top institutions are defined by whether they are ranked among the top 500 research institutions. The data is limited to R1 (very high research activity) and R2 (high research activity) universities.



By contrast, little access to powerful HPC resources is found in states like South Dakota and Utah that have few leading research academic institutions that have the capacity to support HPC systems.

What is important though, is that all regions that lack access to HPC are not the same. Some are doing more AI research than others. For example, while Utah's academic supercomputers are neither particularly large nor particularly powerful, the state is home to the Scientific Computing and Imaging (SCI) Institute, a research institute that focuses on conducting application-driven research in new scientific computing and visualization techniques and tools. The SCI Institute's faculty and alumni are recognized around the world for their contributions to scientific computing and research. South Dakota also has few HPC systems. But unlike Utah, South Dakota has no research facilities identified as conducting high-level research in any field.

The point is, there should be demonstrable evidence that providing access to AI compute in a community, institution, or region will result in an increase in AI research because as explained earlier, democratization is a means to an end, not an end in itself. In cases wherein HPC availability and AI research is low, the Task Force should consider requiring institutions to first increase funding for AI research, prove that they have sought partnerships with industry, or that increasing resources will support AI education and training for underrepresented groups, because there is a risk that investments in AI compute may not return increases in AI research.

We acknowledge that this map is limited because it only shows demand for a subset of academic researchers, not for all researchers. However, as several individuals in the Task Force's workshops have noted, there is little literature on what level and type of compute AI researchers need. Our report offers a starting point, but the Task Force should seek to work with federal agencies and private sector companies, where possible, to obtain additional data on HPC supply and demand.

[C] A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources.

Governance, oversight of strategic direction, and the allocation of federal funds should be made centrally by the private, non-profit corporation. It could have a board of directors appointed by the President of the United States which, after confirmation by the Senate, could serve a fixed length term. In turn, the board would appoint roles for leadership of the corporation, such as president and chief executive officer.



The National Research Cloud could have its own board of directors that oversees day-to-day operations and manages the NRC budget, which could be elected by its member institutions and organizations. Similarly, if the Task Force decides to coordinate the sharing of on-premises systems, this could be governed by a collaborative partnership of participating institutions just as XSEDE is.⁸

[F] *An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls.*

The NAIRR will have a number of distinctive attributes that will make its security somewhat distinct from general-purpose computing architecture.

First, the primary purpose of the NAIRR is to provide researchers with access to advanced, high-performance computing systems, and obtaining time on these systems will be highly valuable. However, stakeholders are likely to disfavor security protocols that impede collaboration or usability.

Second, because the NAIRR may bring together disparate computing systems and distributed data with varying levels of reliability and provenance, there is a risk that the responsibility of cybersecurity will be left to institutions, resulting in a patchwork of security protocols across the country. At the same time, computer security is context- and mission-dependent. A security mechanism designed to enforce a particular policy considered essential for security by one site might unnecessarily block legitimate users of another site.

Fortunately, there are several security solutions that can overcome these constraints and several groups have been thinking about them for a long time. In a 2021 paper titled “Trustworthy Scientific Computing,” Sean Peisert, who leads computer security R&D at Lawrence Berkeley National Laboratory, proposed a model called hardware-based trusted execution environments (TEEs). As Peisert explains, TEEs increase HPC security at a minimal cost to performance by isolating computation and “preventing even system administrators of the machine in which the computation is running from observing the computation or data being used or generated in the computation.”⁹ This paper is part of a larger project that Peisert leads at the Berkeley Lab Computational Research Division, a national laboratory operated by the University of California, to take a broad look at several aspects of security and scientific integrity issues in HPC systems. Since this project has already begun testing and identifying the security requirements of national HPC resources, the Task Force should

⁸ “XSEDE Governance,” last updated May 17, 2021, <https://www.xsede.org/about/governance>.

⁹ Sean Peisert, “Trustworthy Scientific Computing,” *Communications of the ACM*, 64(5), (May 2021), DOI: 10.1145/3457191.



seek to work with this group and others like it to get a fuller understanding of what the security requirements of the NAIRR will likely involve.

2. Which capabilities and services provided through the NAIRR should be prioritized?

The Task Force should prioritize the development of a service-oriented architecture, which would integrate widely divergent components in the NAIRR by providing users with a common interface and a set of standard protocols for them to efficiently access the tools they need.

On one hand, the distributed framework we have proposed for the NAIRR offers an operating model that is flexible enough to adapt to new scenarios, resources, and computing capabilities. Resource diversity is important to ensure AI researchers can remain competitive. Indeed, research and advisory firm Gartner predicts that by 2025, “traditional computing technologies will hit a digital wall, forcing the shift to new computing paradigms such as neuromorphic computing.”¹⁰ There is also already a growing market for emerging AI chips that are specialized to best support different AI capabilities and services. For instance, field programmable gate arrays (FPGAs), which are AI chips mostly used to apply trained AI algorithms to new data inputs, and application-specific integrated circuits (ASICs), which can be used for either training or inference tasks, have seen considerable adoption recently.¹¹

However, a single resource made up of heterogeneous computing systems and data with different architectures, interconnects, memory, and authentication policies presents practical challenges to researchers trying to execute services on the NAIRR and technical developers of the NAIRR who will need to create portals, gateways, and workflow engines for it. Fortunately, many of these problems are not new—just more difficult to solve at scale. XSEDE presents a promising example of how to enhance interoperability and cross-platform functionality. As a “single virtual system that scientists can use to interactively share computing resources, data, and expertise,” XSEDE uses service-oriented architecture to guide users through the different services and capabilities NSF’s resources offer, enabling them to efficiently access their desired functionality.

A tougher problem with heterogeneous computing systems that was raised in a recent Task Force workshop is that it will be more difficult for very computationally intensive problems to be executed because users will have to deal with load balancing over different systems,

¹⁰ Kasey Panetta, “Gartner Top 10 Strategic Predictions for 2021 and Beyond,” *Gartner blog*, October 21, 2020, <https://www.gartner.com/smarterwithgartner/gartner-top-10-strategic-predictions-for-2021-and-beyond>.

¹¹ Saif M. Khan, “AI Chips: What They Are and Why They Matter,” (CSET, April 2020), <https://cset.georgetown.edu/publication/ai-chips-what-they-are-and-why-they-matter/>.



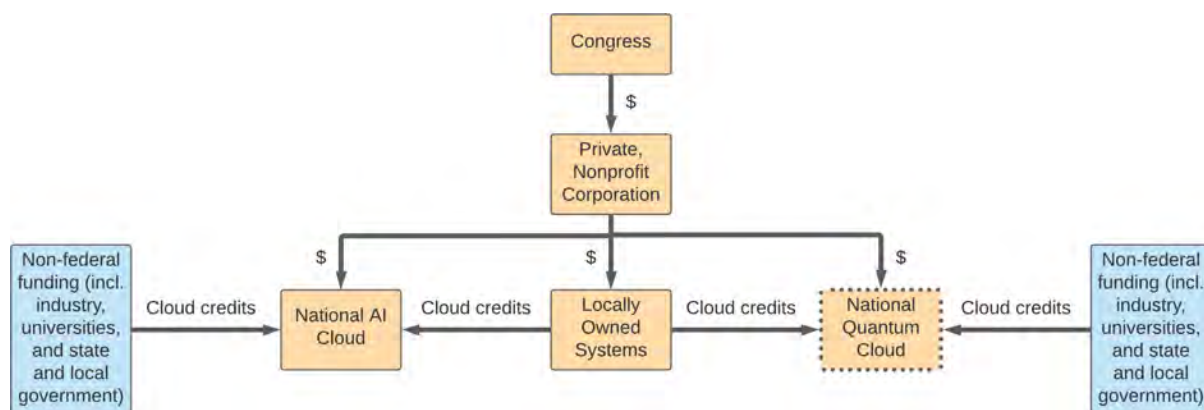
interoperability, resource selection, among other challenges. However, the Task Force should consider that the ultimate goal of the NAIRR is to democratize access to spur AI innovation, bolster U.S. competitiveness in AI, and bridge the “compute divide.” The “long tail” of AI researchers that have more modest computational needs represent, in aggregate, the majority of AI researchers and a significant portion of AI advances. The NAIRR should therefore prioritize capabilities and services that meet the majority of AI researcher needs.

In the long run, a distributed framework could also enable the NAIRR to expand to include different technologies. Most importantly, the Task Force should consider in its roadmap how such a resource could incorporate resources for quantum computing. Because quantum computers are very specialized and expensive to develop, few universities provide access to these systems to support research activities. Instead, most academic researchers access these systems through quantum clouds—services that provide remote access to quantum systems through existing Internet infrastructure. Companies such as Amazon and Microsoft have already begun to make access to quantum computers available through their quantum computing-as-a-service (QCaaS) offerings, which are fully managed services that enable researchers and developers to begin experimenting with systems from multiple quantum hardware providers in a single place. Even with declining computing costs though, the costs and know-how for using advanced computing, including QCaaS solutions, will remain out of reach for many academic researchers.¹²

While AI and quantum computing differ, the crux of the problem is the same: How can the United States provide academic researchers with affordable access to high-end computing resources in a secure environment? Rather than reinventing the wheel, Figure 3 below illustrates how the scope of the NAIRR could be adapted to include additional resources to support quantum computing research.

¹² Hodan Omaar, “The Case for a National Quantum Computing Research Task Force in the United States,” June 9, 2021, <https://datainnovation.org/2021/06/the-case-for-a-national-quantum-computing-research-task-force-in-the-united-states/>.

Figure 3: Expanding the scope of the NAIRR to include quantum computing.



3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

The ethical considerations regarding the NAIRR fall into two main buckets. One is how to ensure the allocation of resources in the NAIRR are fair and the other is how to ensure those resources are used to advance ethical and responsible AI research.

The first question is essentially a cake cutting problem, which is the challenge of allocating a single divisible, continuous, resource in a fair and equitable manner.¹³ The “cake” in this case is the NAIRR and individuals have different preferences over different pieces (because they will be pursuing different types of research and different systems within the NAIRR will be better suited to their needs). How should one split the cake so that it is fair in the sense of distributional fairness, understood as maximizing everyone’s utility, and in the sense of not having disparate impact across protected groups?

Such practical problems are studied in mechanism design, a field of economics that studies the mechanisms through which a particular outcome or result can be achieved. Mechanism design can help bring analytic clarity to policy goals. Consider the statement: “The NAIRR should provide AI compute to the greatest number of AI researchers.” This directive could be operationalized multiple ways. One way would be to minimize the total number of AI researchers that have less than some threshold of AI compute. Another way would be to first

¹³ Rediet Abebe et al., “Fair Division via Social Comparison,” *AAMAS '17: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, (May 2017), pages 281–289, DOI: 10.5555/3091125.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Center for Democracy and Technology

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

October 1, 2021

Office of Science and Technology Policy and the National Science Foundation
Attn: Wendy Wigen, NCO,
2415 Eisenhower Avenue,
Alexandria, VA 22314, USA

Re: Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource

(submitted by email)

In response to the National Science Foundation and the Science and Technology Policy Office's request for information on an implementation plan for an National Artificial Intelligence Research Resource, the Center for Democracy & Technology would like to offer the following comments:

1. What options should the Task Force consider for any of the roadmap elements A through I above, and why?

B.ii: A governance structure for the Research Resource, including oversight and decision-making authorities.

In designing a governance structure for the Research Resource, the NAIRR Task Force should first consider the types of decisions that decision-making authorities will be called upon to make, and establish a governance structure that will involve individuals with sufficient subject matter expertise in these topics and that will ensure transparency, accountability, and fairness. Among other issues, decision-makers will have to determine:

1. Who should be granted access to the Research Resource, what criteria AI researchers and students must satisfy to gain access to the Research Resource, and how researchers and students will be vetted for compliance with that criteria.
2. What limits will be placed on researcher and students' use of the Research Resource once they are granted access. These limits should be designed to protect individual privacy and prevent unethical uses of the Research Resource.

1401 K Street NW, Suite 200 Washington, DC 20005

3. How the limits on use of the Research Resource will be enforced, including how violations will be reported or otherwise detected, investigated, and found to be substantiated or unsubstantiated, and the penalties for substantiated violations of the limits.
4. How the NAIRR can empower researchers and students to make the best use of the Research Resource and ethically use the Research Resource, through technical support, training, and other resources.
5. What data should be included in the Research Resource, what criteria data must satisfy to gain access to the Research Resource—including criteria concerning lack of bias, privacy, and intellectual property rights—and how data will be vetted for compliance with that criteria.
6. What data included in the Research Resource must be kept confidential and accessible only to vetted AI researchers and students, and what data can be made publicly available.

These considerations should guide the development of a diverse and sufficiently empowered governance structure for the NAIRR. That structure should be transparent in terms of who will make decisions, the process for doing so, and how those decisions will be communicated to the public and/or involved parties. The structure should be fair: for example, a researcher that is denied access to the resource should have a means to appeal or contest that denial. And the governance structure should make clear who is accountable for the various decisions that will need to be made.

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.

On the provision of curated data sets, the NAIRR Task Force should provision and disseminate curated data sets in ways that allow researchers to reproduce and build upon existing research. Not being able to reproduce research results, or what scientists call the “replication crisis”, is a significant challenge for science overall, and one of its causes is insufficient access to the underlying data researchers use. One approach to address this is through the use of open data — data that is freely available to use and redistribute with few or no restrictions. Requirements for certain government agencies to open certain data sets already exist at the federal,¹ state, and local levels. However, given that the NAIRR will most likely include collaborations across different research groups from the public and private sectors, an

¹ DCAT-US Schema v1.1 (Project Open Data Metadata Schema), <https://www.data.gov/>

open data approach is important. In practical terms this means providing data in open formats (i.e., not dependent on particular software), with little or no restrictions (e.g., registration requirements). It also means publishing metadata,² creating public APIs where feasible, and being clear about how the data is licensed (ideally licensing similar to that of Creative Commons, or when possible, placing the data in the public domain).

While open data can help address the reproducibility challenge, access alone is not enough.³ The NAIRR should also develop guidelines on how to create documentation to accompany each dataset that will include how the dataset was constructed (e.g., labelling, calculation of new variables, etc.) and how it was used in the research. These are not common practices for researchers, and so the NAIRR Task Force should also explore ways to incentivize and guide researchers through these steps.

The availability of open data can also promote greater equity in the access and use of the NAIRR, particularly among researchers with limited resources or those outside of research networks of scholars willing to share their data. However, open data practices alone cannot address the equity problem, and the NAIRR will have to be intentional in understanding and addressing the data needs of researchers in non-traditional research organizations (e.g., journalists, civil society) and academic groups that have traditionally collaborated less directly with entities such as the NAIRR itself.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research.

CDT recommends that the NAIRR Task Force conduct a Human Rights Impact Assessment (HRIA) of the Research Resource and build future HRIAs into the governance structure of the Research Resource. HRIAs are a method of reviewing and monitoring particular projects or activities to offer “guidance and practical tools” using a human rights-based approach.⁴ While businesses may undertake HRIA to assess consistency of their activities with the United Nations Guiding Principles on Business and Human Rights, governments can and should also engage in HRIAs to analyze the consistency of their activities

² See for example DCAT-US Schema v1.1 (Project Open Data Metadata Schema), <https://resources.data.gov/resources/dcat-us/>

³ See for example Hardwicke, T. E., Mathur, M. B., MacDonald, K., Nilsson, G., Banks, G. C., Kidwell, M. C., ... & Frank, M. C. (2018). Data availability, reusability, and analytic reproducibility: Evaluating the impact of a mandatory open data policy at the journal *Cognition*. *Royal Society open science*, 5(8), 180448.

⁴ The Danish Institute for Human Rights, *Welcome and Introduction, Human Rights Impact Assessment Guidance and Toolbox* (2020) at 4, https://www.humanrights.dk/sites/humanrights.dk/files/media/document/DIHR%20HRIA%20Toolbox_Welcome_and_Introduction_ENG_2020.pdf.

with international human rights principles. The NAIRR Task Force should engage in a HRIA of the proposed Research Resource and incorporate learnings from that assessment into its design of the governance structure for the Research Resource. In addition, CDT recommends building ongoing and periodic HRIAs of the Research Resource into its governance structure, with a particular emphasis on analysis of whether and how the Research Resource is impacting individual privacy or enabling biased, discriminatory, inequitable, or unethical research or application of AI.

In designing and implementing these HRIAs, the NAIRR Task Force may wish to consider the Human Rights Impact Assessment Guidance and Toolbox from the Danish Institute for Human Rights.⁵ In addition, the NAIRR Task Force could look to HRIAs conducted of other data-sharing endeavors, such as the HRIA of the Global Internet Taskforce for Combatting Terrorism, through which member companies make use of a hash-sharing database and URL sharing to identify and screen user-generated terrorist content.⁶

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

We commend the NAIRR Task Force for considering issues of equity, fairness, bias, civil rights, transparency, and accountability during the roadmapping and planning process for the NAIRR. These are challenging and complex issues that should be considered early on in any development process for AI-based frameworks and resources and, as noted above, should be addressed as part of the NAIRR's governance structure.

CDT would like to offer a few suggestions aimed at helping to ensure that issues of equity and accountability are integrated into the NAIRR infrastructure.

1. The NAIRR Task Force should start with its own research into privacy and equity harms from prevalent AI systems. This includes regularly consulting with a range of individuals and

⁵ See Danish Institute for Human Rights, *Human Rights Impact Assessment Guidance and Toolbox*, above.

⁶ Global Internet Forum to Counter Terrorism, *Human Rights Assessment*, https://gifct.org/wp-content/uploads/2021/07/BSR_GIFCT_HRIA.pdf.

organizations with personal experience and technical and policy expertise, from affected consumers and civil society groups to academics and AI researchers. Consumers and civil society groups can speak to the direct impacts of AI systems' discriminatory outcomes, offering valuable insight into the information asymmetry involved in AI systems, the burden that being researched can pose to consumers, and the benefit that direct access to resources like the NAIRR would provide in disputing unfair AI-driven outcomes. Academics and AI researchers can speak to commonly researched types of AI systems, underlying AI issues and goals driving the research, research methodologies, and the ultimate effectiveness of selected data sets and methodologies. This discourse can shape the Task Force's ability to anticipate how the NAIRR's substance and degree of accessibility might contribute to reducing or preventing inequities and privacy violations. The Task Force should build these considerations into its privacy and civil rights and civil liberties requirements, as well as a process to modify the requirements as soon as they are found to be inadequate.

2. The NAIRR infrastructure should include dedicated resources for projects that enable factfinding and research about equity and bias in existing AI systems and methods. There are numerous examples of biased AI-based systems causing harm to people,⁷ and the NAIRR should offer priority to researchers and projects that seek to uncover, understand, and, where possible, correct these equity issues.
3. There are many cases where AI systems produce biased outcomes that stem, at least in part, from biased training datasets.⁸ The NAIRR should audit any data sets it provides for bias. The NAIRR should seek to build and provide unbiased datasets where possible, while understanding that mitigating bias in datasets is complex and, if done incorrectly, may itself introduce bias concerns. For instance, consider a dataset of test scores, where students of color have disproportionately worse scores due to bias in the test design. It may seem that a solution

⁷ Meredith Broussard, When Algorithms Give Real Students Imaginary Grades, N.Y. Times (Sept. 8, 2020), <https://www.nytimes.com/2020/09/08/opinion/international-baccalaureate-algorithm-grades.html>; A-levels and GCSEs: How Did the Exam Algorithm Work?, BBC News (Aug. 20, 2020), <https://www.bbc.com/news/explainers-53807730>; Rebecca Koenig, Can Algorithms Select Students "Most Likely to Succeed"?, Slate (July 10, 2020), <https://slate.com/technology/2020/07/college-admissions-algorithms-applications.html>; Shirin Ghaffary, The Algorithms that Detect Hate Speech Online Are Biased Against Black People, Vox (Aug. 15, 2019), <https://www.vox.com/recode/2019/8/15/20806384/social-media-hate-speech-bias-black-african-american-facebook-twitter>; Karen Hao, The Coming War on The Hidden Algorithms That Trap People in Poverty, MIT Tech. Review (Dec. 4, 2020), <https://www.technologyreview.com/2020/12/04/1013068/algorithms-create-a-poverty-trap-lawyers-fight-back/>.

⁸ Jeff Larson et al., How We Analyzed the COMPAS Recidivism Algorithm, (May 23, 2016), <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>.

would be to remove low scores from students of color to avoid entrenching the original biased test design in the resulting score dataset. However, that dataset would now have more limited information about students of color than about white students, which is itself a form of bias. Consequently, any attempt to correct bias in datasets must be an iterative and rigorous process that avoids introducing new forms of bias.

In many cases, collecting unbiased data may not be possible: bias in datasets typically stems from bias in the environment where the data is collected, meaning that existing and pervasive societal biases will inevitably be reflected in real-world data. In such cases, the NAIRR should identify the purpose for developing each data set and using the types of data involved, the rationale for relying on particular sources of data, and the issues on which each data set provides relevant insights. If these elements are not readily articulable, this could signal privacy and equity risks that need to be more deeply explored, in which case the NAIRR either should not include the dataset in question or should provide clear information about the bias(es) in the dataset so users of the dataset are aware of and can appropriately account for such bias.⁹

The NAIRR Task Force should take advantage of existing efforts within the federal government to understand and limit bias and inequity in AI-based systems. For example, the NAIRR should work in concert with efforts by the National Institute of Standards and Technology (NIST) to improve explainability in AI¹⁰ and understand and mitigate bias in AI.¹¹

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The NAIRR's commitment to offering training to bring more people into the AI R&D community will be an important element of democratizing the field, but in order to be effective, that training will need to take into account the existing inequities in the AI community, and the tech sector more broadly. For instance, trainings or resources targeted at university computer science students to encourage them to

⁹ Timnit Gebru et al., Datasheets for Datasets, (Mar. 19, 2020), <https://arxiv.org/abs/1803.09010>.

¹⁰ P. J. Phillips et al., Four Principles of Explainable Artificial Intelligence (Draft), (Aug. 18, 2020), <https://www.nist.gov/publications/four-principles-explainable-artificial-intelligence-draft>.

¹¹ NIST Proposes Approach for Reducing Risk of Bias in Artificial Intelligence, (June 22, 2021), <https://www.nist.gov/news-events/news/2021/06/nist-proposes-approach-reducing-risk-bias-artificial-intelligence>.

specialize in AI will not address the lack of diversity amongst computer science students at the university level, and thus will be limited as far as democratizing and diversifying the field. The NAIRR can help to overcome these limitations by acknowledging and working to combat existing challenges. For example, this could mean offering trainings and resources designed to pull students and the general public into AI from a broader variety of fields than just computer science, including those with a higher concentration of populations that are currently underrepresented in the AI field, such as nursing or teaching.

The Task Force should also consider the need for a clear process allowing equitable access to the NAIRR and to resources within the NAIRR that help contextualize and explain data sets for consumers' understanding. Providing avenues for engagement and participation from civil society and advocacy groups that work on behalf of marginalized communities (rather than just academic or industry AI researchers) is critical to avoiding discriminatory outcomes and enabling those communities to self-advocate.¹² This is particularly important given that marginalized people experience greater barriers to entry and advancement within the AI research sector.¹³ Without their own access to resources like the NAIRR, marginalized communities would have to depend on qualified researchers who may not prioritize the AI-related issues that concern and harm consumers most.¹⁴ In certain situations, researchers' analysis may also not be easily available to the public due to other interests involved in the research endeavor.¹⁵ Overall, the Task Force must actively examine its processes for building, curating, and providing access to AI resources to determine whether their use remains appropriate and continues to serve the public interest.

Overall, CDT commends the Office of Science and Technology Policy and National Science Foundation for undertaking the development of a resource as complex and potentially valuable as the NAIRR, and for considering important questions of equity, privacy, and democratic access from the beginning.

¹² Kathryn L.S. Pettit et al., Urban Institute, Putting Open Data to Work for Communities (2014), available at <https://www.urban.org/sites/default/files/publication/22666/413153-Putting-Open-Data-to-Work-for-Communities.PDF>.

¹³ Ebony O. McGee, *Let's Remake Racially Unsafe STEM Educational Spaces*, Higher Education Today (Feb. 11, 2021), <https://www.higheredtoday.org/2021/02/11/lets-remake-racially-unsafe-stem-educational-spaces/>.

¹⁴ Florence Ashley, Accounting for Research Fatigue in Research Ethics, 35 *Bioethics* 270, 272 (2021), available at https://www.florenceashley.com/uploads/1/2/4/4/124439164/ashley_accounting_for_research_fatigue_in_research_ethics.pdf.

¹⁵ Brian Resnick and Julia Belluz, *The War to Free Science*, Vox (July 10, 2019), <https://www.vox.com/the-highlight/2019/6/3/18271538/open-access-elsevier-california-sci-hub-academic-paywalls>.



Sincerely,

Hannah Quay-de la Vallee, *Senior
Technologist, CDT*

Gabriel Nicholas
Research Fellow, CDT

Ridhi Shetty
Policy Counsel, Privacy & Data Project, CDT

Dhanaraj Thakur
Research Director, CDT

Caitlin Vogus
Deputy Director, Free Expression Project, CDT

1401 K Street NW, Suite 200 Washington, DC 20005

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Cerner Corporation

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

September 20, 2021

Attn: Wendy Wigen
National Coordination Office (NCO)
2415 Eisenhower Avenue
Alexandria, Virginia 22314, US

RE: [Request for Information \(RFI\) on an Implementation Plan for a National Artificial Intelligence Research Resource](#)

Dear Ms. Wigen,

Cerner Corporation appreciates the opportunity to submit public comment on the White House Office of Science and Technology Policy and National Science Foundation's RFI on an Implementation Plan for a National Artificial Intelligence Research Resource. As a leading supplier of clinical and management information systems we believe our experience provides us with valuable insight in this subject area and are grateful for the ability to share that insight.

If you have any questions or if we can provide any additional information, please do not hesitate to contact me at [REDACTED]

Sincerely,

John Travis
Vice President & Regulatory Strategy Executive
Cerner Corporation

First, Cerner would like to express our appreciation to the Office of Science and Technology Policy and National Science Foundation, and the National Artificial Intelligence Research Resource (NAIRR) Task Force, for their focus on this subject. We believe it is a critical topic for all industries – particularly as Artificial Intelligence (AI) and Machine Learning (ML) continue to become more established parts of research, data analytics, and software development. We believe the NAIRR Task Force’s focus and investment with this Request for Information (RFI) and subsequent review is a critical step towards achieving a higher level of focus on AI and ML within the industries.

This RFI relating to AI research spans all industries. Cerner’s responses to the RFI questions are from a healthcare industry perspective.

As various types of data contribute to a national AI database, it is critical to consider the sensitive nature of healthcare data, whether it be clinical or financial, compared to many other types of data. People may fear discrimination from an employer or from an insurance company based on disclosure or secondary use of certain health information. Needless to mention, the patient safety considerations and potential discriminatory uses of health information is heightened just by the mere fact that it is data related to human health.

It is for these reasons that the data (and NAIRR as a whole) should be segmented based on purpose and scope within the infrastructure. For example, for the reasons outlined above, healthcare data should be treated differently than other types of data. Access to healthcare data must adhere to Health Insurance Portability and Accountability Act (HIPAA) Privacy and Security rules, including de-identification via removal of Personal Health Information (PHI). If raw PHI is needed to compute a derived value, when the computation is completed, the PHI must be discarded.

1. WHAT OPTIONS SHOULD THE TASK FORCE CONSIDER FOR ANY OF ROADMAP ELEMENTS A THROUGH I BELOW, AND WHY? [PLEASE TAKE CARE TO ANNOTATE YOUR RESPONSES TO THIS QUESTION BY INDICATING THE LETTER(S) OF THE ITEM (A THROUGH I IN THE LIST BELOW) FOR WHICH YOU ARE IDENTIFYING OPTIONS.]

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and

Cerner comments

Unless AI ownership is segmented out in different categories by type and purpose of collection (e.g., healthcare, defense, etc.), it will likely be infeasible to have a single resource that can be fully shared and serve all unique needs.

September 20, 2021

Page 3

ii. A governance structure for the Research Resource, including oversight and decision-making authorities

Cerner comments

It is critical that the governance structure for the NAIRR involve representation from knowledgeable and relevant public-private entities for the types of data and purposes of collection that prevail. This is necessary to ensure the requisite knowledge set is present to appropriately inform key decisions about the data, the development lifecycle of that data, AI/ML techniques, the implications of obtaining that data, and how the data should be handled related to privacy.

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure

Cerner comments

Healthcare poses a unique use-case which necessitates special considerations such as elevated protections for PHI. The associated infrastructure strategy for the NAIRR must account for that.

2. WHICH CAPABILITIES AND SERVICES (SEE, FOR EXAMPLE, ITEM D ABOVE) PROVIDED THROUGH THE NAIRR SHOULD BE PRIORITIZED?

Cerner comments

Priority needs to be on publishing standards as part of this resource that provide the industry with guidelines for development and testing of AI software. As the FDA develops these guidelines for health information through ongoing guidance development for AI/ML, this is an opportunity to integrate with that organization to establish consistency between what they're providing and how this resource functions.

3. HOW CAN THE NAIRR AND ITS COMPONENTS REINFORCE PRINCIPLES OF ETHICAL AND RESPONSIBLE RESEARCH AND DEVELOPMENT OF AI, SUCH AS THOSE CONCERNING ISSUES OF RACIAL AND GENDER EQUITY, FAIRNESS, BIAS, CIVIL RIGHTS, TRANSPARENCY, AND ACCOUNTABILITY?

Cerner comments

While it is important to cultivate innovation, the principles of ethical and responsible research and development of AI cannot be neglected. The code of conduct developed to provide guidance on how to abide by these principles should encourage self-regulation and self-reporting. These principles fall into the following three categories:

- Data Responsibility
- Ethical Principles

September 20, 2021

Page 4

- Transparency and Accountability

Data Responsibility

Data policies and procedures should be documented, maintained, and enforced to:

- Ensure technical security and user security
- Provide system-level logging and security
- Provide incident management
- Provide change management
- Provide contingency planning for disaster recovery and resiliency
- Not use or disclose/share the data in ways other than stated, or as otherwise required by law
- Safeguard the data to prevent such misuse or unauthorized disclosures
Report any misuse or unauthorized disclosure as soon as known
- Refrain from identifying or contacting subjects represented by or within the data

Ethical Principles

Similar to principles developed by organizations like the [Center for Practical Bioethics](#), the [Berkman Klein Center at Harvard University](#) the [Stanford Institute for Human-Centered Artificial Intelligence](#), ethical principles must be applied to the use of the data.

- Literature research should be performed on the subject of interest with special attention on the topics of bias and fairness.
- Communication with Subject Matter Experts and other stakeholders is critical.
- Data sources must be understood and collection methods to identify potential opportunities for data bias must be examined.
- Diverse cohorts should be used in order to mitigate bias.
- Possible operational and usage patterns should be defined to identify any potential consequences of the AI's use that may affect some sensitive groups (e.g., racial, and ethnic minorities).
- Possible AI design approaches that may minimize bias should be identified.
- Fairness and bias should be measured using quantifiable, mathematical metrics.
- Expectations for transparency to users from the AI should be defined.
- Transparency of the AI development and validation process to consumers of the AI as it relates to potential for bias should be ensured.
- Analysis of appropriate performance and bias metrics for relevant subgroups should be performed. Typical subgroups include gender, race, ethnicity, age, marital status, location and time period.
- When a subgroup exhibits unacceptable performance that cannot be improved, the AI should be documented as not to be used on that particular subgroup.

September 20, 2021

Page 5

Transparency and Accountability

As AI is used to drive toward positive outcomes, the most important aspect is to ensure the algorithms are monitored over time to retain quality and accuracy.

Quality management systems should be in place to ensure processes are followed, information is used in the same way, and methods are consistent. This is not to say a national quality regulatory framework that requires submission thru a regulatory body or something along the lines of a 510(k) is needed, rather that quality management systems should be utilized as they exist for software today, similar to the [FDA's Quality System Regulation \(QSR\)](#) and the [International Organization for Standardization's \(ISO\) quality management system](#).

If the monitoring identifies a need for change, end users must be notified, use of the intelligence must be discontinued if warranted until a change can be implemented and note of corrective actions must be taken. If the reason a change is needed is to prevent an adverse event or significant financial, regulatory or privacy burden, a process should be in place to determine root cause and create mitigation(s) to prevent such an issue from reoccurring.

Adverse events should be detected during monitoring by:

- Applying the AI to two different demographics. Comparing the outcomes against both demographics can proactively assess how much risk is due to the demographics.
- Examining the AI results for an unexplained minority group effect that is statistically significant compared to the median.
- Explicitly creating a quality measure of disparity that can be monitored over time. Changes in this quality measure can be investigated as new data is introduced.

Bias should be combated through responsible development practices and monitoring of AI in real world use. A monitoring framework and ecosystem should exist with the following capabilities to:

- Detect a worsening outcome disparity or performance bias. Mitigation options should include:
 - stopping use of the AI where it is performing poorly,
 - modifying the AI to correct the disparity or bias, and
 - proactively seeking out and considering users' feedback on their perception of bias.
- Regularly produce automated reports on the distribution of data inputs, performance and bias so they can be reviewed for:
 - changes in the mean or variance of a continuous data element,
 - changes in patterns of missing data,
 - changes in the distribution of categorical data – especially if a new category appears, as well as,

September 20, 2021

Page 6

- changes in results based on demographics factors.
- Produce near real-time alerting when a critical failure occurs.
- Proactively look for measures that may indicate a disparity in outcomes arising from the use of AI across demographic lines such as race, age, gender, or some combination of them.

4. WHAT BUILDING BLOCKS ALREADY EXIST FOR THE NAIRR, IN TERMS OF GOVERNMENT, ACADEMIC, OR PRIVATE-SECTOR ACTIVITIES, RESOURCES, AND SERVICES?

Cerner comments

No comments

5. WHAT ROLE SHOULD PUBLIC-PRIVATE PARTNERSHIPS PLAY IN THE NAIRR? WHAT EXEMPLARS COULD BE USED AS A MODEL?

Cerner comments

As mentioned above in question #1 under letter item B.ii, a governance board with representation from public-private participants is critical to contributing to oversight and decision-making about the data, understanding the data lifecycle, and protecting the privacy of the data (e.g., PHI, etc.) so that implications of decisions are well understood.

Accordingly, a distinction or separation should also exist that would delineate the management of the data assets from public policies/laws and the actual AI development. Public policies and laws are government core competencies. However, administration of the data asset and associated technical aspects may be more suited for a public-private partner.

6. WHERE DO YOU SEE LIMITATIONS IN THE ABILITY OF THE NAIRR TO DEMOCRATIZE ACCESS TO AI R&D? AND HOW COULD THESE LIMITATIONS BE OVERCOME?

Cerner comments

No comments

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Sean Ekins

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource

Responses from Dr. Sean Ekins, CEO, Collaborations Pharmaceuticals, Inc.

question 1

A. yes - set some goals, preferably attainable ones that will matter for the nation and not get undercut every 4 years if the government changes.

B. Yes - I would have thought about this and decided upon it before starting.

C. Yes - you should have done that already - the last thing you need are 20-30 organizations arguing over it.

D. Yes – I would spend the most time and effort on this - This is the reason a national AI resource should exist.

E. Yes - if the data is funded by government, it should be open full stop - there should be no barriers to data, if there are, defund the agency.

F. No - This should absolutely be done independently.

G. No - again this should be assessed by an independent agency.

H. No - if you are going to set something up like this then it is going to need a decades long horizon..This is not something you fund for a few years and then expect to be self funding if this is going to be so critical to the nation's infrastructure. The last thing this needs is private funding or partnerships. The last thing this needs is a Facebook, Google, Microsoft or others involved - This is for the nation and its long term outlook and not for some companies to direct and profit off.

I. No - It has a reason to exist I would not spend effort in justifying existence.

question 2

D

Question 3

For one make sure that from the outset the people involved with setting it up are actually not all old white men.

Question 4

You do not have to look far, EPA, NIH are great places to start. I am pretty sure that if there is money to be spread around you will get enough academics and companies running to have a seat at the table.

Question 5

None. No partnerships. Create the resources so anyone can access data and hardware. Companies should have no preferred access any more so than any other individual or organization in the country.

Question 6

As soon as you get companies involved it will be seen as tainted and not democratic. if anyone can use and access the facilities it should be a level playing field, Ivy league schools do not get priority over other schools, billion-dollar companies do not get any more priority than a small start-up and the individual citizen has to have as much right to access as anyone from a large organization.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Computing Community Consortium, Computing Research Association- Industry, Association for the Advancement of Artificial Intelligence

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Computing Community Consortium (CCC), Computing Research Association-Industry (CRA-I), and The Association for the Advancement of Artificial Intelligence (AAAI) Responds to the Implementation Plan for a National Artificial Intelligence Research Resource

September 2021

Liz Bradley (CCC Chair and University of Colorado Boulder), Nadya Bliss (Arizona State University), William Gropp (University of Illinois Urbana-Champaign), Helen Nissenbaum (Cornell Tech), Chris Ramming (CRA-I and VMware), Ann Schwartz (CCC), Bart Selman (AAAI President and Cornell University), Helen Wright (CCC)

Response to Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource:

<https://www.federalregister.gov/documents/2021/07/23/2021-15660/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>

The following is a joint response from the [Computing Community Consortium](#) (CCC), [Computing Research Association-Industry](#) (CRA-I), and the [Association for the Advancement of Artificial Intelligence](#) (AAAI). We offer this joint response to the following points from the RFI drawing from the extensive discussions within the national AI research community that arose while developing [A 20-Year Community Roadmap for Artificial Intelligence Research in the US](#) (AI Roadmap) as well as conversations among the newly formed CRA-I Steering Committee and community. All text which is bold and italicized is directly from the RFI. Text which is a regular font is the response.

- 1. *What options should the Task Force consider for any of the roadmap elements A through I, and why?***
 - a. *Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;***

A NAIRR should support work on a range of core challenges in AI research and development. In our AI Roadmap, we identified three core themes central to the further development of AI:

- i. Integrated intelligence, including developing foundational principles for combining modular AI capabilities and skills, approaches for contextualizing general capabilities to suit specific uses, creation of open shared repositories of machine understandable world

knowledge, and understanding human intelligence both to inspire novel AI approaches and to develop models of human cognition.

- ii. Meaningful interaction, comprising techniques for productive collaboration in mixed teams of humans and machines, combining diverse communication modalities (verbal, visual, emotional) while respecting privacy, responsible and trustworthy behaviors that can be corrected directly by users, and fruitful online and real-world interaction among humans and AI systems.
- iii. Self-aware learning, developing robust and trustworthy learning, quantifying uncertainty and durability, learning from small amounts of data and through instruction, incorporating prior knowledge into learning, developing causal and steerable models from numerical data and observations, and learning real-time behaviors for intentional sensing and acting.

Each of these areas requires a setting with interactive data collection, where machines and humans interact and machine learning is a dynamic process involving active learning in a changing environment. Static datasets can only provide a starting point for tackling these research challenges. Data collection and generation needs to be dynamic, guided by AI decision making itself. Examples are (1) the use of AI for scientific discovery, where AI methods guide the scientific experimentation in a continuous loop of data analysis followed by subsequent experimentation and new data collection and analysis, (2) work on human-robot interaction where the AI system uses active learning to explore and map its environment and dynamic interactions with users, (3) in AI for healthcare, where AI data analysis and decision support systems are part of the clinical process, continuously monitoring health data and providing continuous guidance, and (4) personalized AI for lifelong learning, where the AI education support system adjusts to new learning goals and needs over time.

The AI Roadmap therefore recommends a broader perspective on the required shared National AI Research Infrastructure. Such infrastructure would embed AI research in actual application environments. One example would be an "AI-ready research hospital" where AI researchers from around the country could work (likely remotely) with clinicians and other medical staff to develop AI for health applications. The AI systems

would provide data analysis and input to the clinical process, while continuously monitoring patient progress and multiple sources of health data. Another example would be a shared materials science research facility that enables the development of AI-driven automated experimentation searching for new materials. A final example would be a shared facility for the development of interactive robots for assisted living. Active learning and continuous data collection are again central to the development of the next generation of assistive robots.

These examples all point to the need for shared AI Research Infrastructure resources that can provide active continuous data collection in dynamic environments to train the next generation of AI decision support systems. The AI systems need to be integrated into these environments to allow for active learning (where the AI system itself decides what data to collect next), and for dynamic control and decision making.

- b. *A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:***
- i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and... ?***
 - ii. A governance structure for the Research Resource, including oversight and decision-making authorities;... ?***

In looking at the organization and management of the NAIRR, it is important to recognize both the key importance of data and the reality that that data will be distributed across the country, in commercial clouds, at national research facilities, and at academic institutions. An NAIRR needs to place adequate computing near the data, which implies that there are multiple sites providing computing resources for NAIRR. This is not a problem, and in fact offers some advantages. Unlike Leadership Computing Facilities, it isn't necessary to have all of the computing in one place, and few single jobs would need to use all of the computing resources at once in a tightly-coupled way – and research that needs such resources could instead use DOE Exascale systems through the DOE INCITE program. An advantage is that with multiple sites, and with both public and private providers, the NAIRR can include a broad menu of different computing technologies and ensure that there are

frequent updates and incentives for exploring innovative new hardware and software.

NSF has a successful model for this approach. NSF funds advanced computing systems, including innovative pilots, and provides a virtual organization through the Extreme Science and Engineering Discovery Environment (XSEDE). The flexibility of this organization was recently demonstrated in the COVID-19 HPC Consortium, where XSEDE played a key role in managing the review of proposals and allocations of the resources, which included contributions from commercial clouds, national laboratories, and academic supercomputer centers. While XSEDE is ending next year after many successful years of operation, NSF is continuing this virtual organization model through a new program, ACCESS. While the details of the service model used in XSEDE and ACCESS will be different for NAIRR, the use of a virtual organization following the XSEDE model provides a tested, successful model for providing access to research computing while ensuring ongoing innovation, and supporting a distributed model essential for providing access to the many disparate data sources. One difference from the XSEDE model is that NAIRR should also have regular, planned investment in both hardware and software, something that is not part of XSEDE (which only provides support services – hardware and software are funded through separate competitive NSF solicitations).

d. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

- i. Create something similar to the National AI Research Centers, which are intended to create unique and stable environments for large multidisciplinary and multi university teams devoted to long-term AI research and its integration with education. A shared computing infrastructure should be:
 1. Funded through decade-long commitments, to provide stability and continuity of research

2. Multi-stakeholder centers with a smaller core set of partners and a large network of affiliated educational institutions, national labs, and industry. Promoting collaboration between academia, industry, and government will enable cross-cutting research and technology transition.
3. Multidisciplinary in the expertise involved in the research areas
4. Multi-faceted in their research goals
5. Effective dissemination vehicles for significant results

e. *An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;*

- i. Access to data sets is a huge problem, and this issue is not limited to government data. The field has matured beyond its initial academic focus on algorithms and theories and into a context of continuous data collection, social and interactive experimentation, and massive amounts of knowledge about a constantly changing world. Building from those foundations, the tech industry has compiled and leveraged massive resources—datasets, knowledge graphs, special-purpose computers, and large cadres of AI engineers—to propel powerful innovations. It is important to create *open* AI platforms and resources, which will be a vast interlinked distributed collection of “AI-ready” resources (such as curated high-quality datasets, software, knowledge repositories, testbeds for personal assistants and robotics environments) contributed by and available to the academic research community, as well as to industry and government.

f. *An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;*

- i. Shared resources (as highlighted in other sections of this response, such as XSEDE and DOE resources) have implemented security requirements and those requirements should be consulted for NAIRR. NAIRR shared resources should as a baseline follow all recommended shared resources requirements of other research infrastructures.
- ii. There are additional considerations in context of AI research around data, results, systems, adversarial AI techniques, and

domain specific security. The following described what is meant by each:

1. Data: well curated data sets from a vast variety of sources are vital to AI research. At the moment, academic researchers do not have access to the same quality of data as industry researchers. A significant benefit of implementation of this resource would be providing access to datasets that the researchers may not have access otherwise. However, different datasets would likely require different restrictions and those will need to be managed through this implementation plan (administratively, legally, technically, etc). Secure multi-party computation, homomorphic encryption, and hardware trusted execution environments (TEEs) can potentially overcome barriers to confidential data sharing and should be considered as long and short-term foundations for the creation and use of relevant datasets
2. Results: given the potential sensitivity of AI algorithms and system results on particular datasets, it will be necessary to make sure that the requirements that are applied to datasets are also applied to the derivative products of those datasets.
3. Systems: as per (1) and (2), a key benefit of the NAIRR would be the ability to seamlessly combine various datasets and results into novel AI systems. The interactions between various components of those systems could potentially introduce either cybersecurity vulnerabilities or dataset biases. Thus, it is important that those are considered as part of access controls and software review. This is another area where potentially NAIRR could learn from other shared resources and/or “app store” implementations (for example, researchers may want to “publish” their algorithms and software modules for other researchers to use, but that “publication” process should include some assessment of potential security vulnerabilities).
4. Adversarial AI: an important area of research is understanding vulnerabilities in AI algorithms. It would be beneficial to the research community if NAIRR could provide a “sandbox” environment for adversarial AI research with appropriate ethical and security restrictions and potential assessment of national security implications.

5. Domain-specific security: given that AI research is performed and can be beneficial to a broad range of application domains, NAIRR should consider if additional security requirements need to be implemented in certain situations. For example, security and access requirements would likely be different between analysis of human genome data and development of AI algorithms for autonomous vehicles. As such, domain experts should be brought in for those situations.

g. *An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;*

- i. For the NAIRR, it is important to remember that all data is sensitive, especially in this AI and Machine Learning environment, all data can implicate all other data.
- ii. Data is a fundamental resource. Due to the enormous interest in all aspects of human and social life, data about people is highly valued. This inevitably raises privacy issues throughout the data cycles of collection, assembly, and analysis. These issues need to be grappled with not as an afterthought, for example, but at the moment of deployment or use, but as an integral consideration at each phase.
- iii. Access to all data needs to be broadened. Whenever you have tasks that are “handed off” to AI systems from a human actor, there also needs to be accountability, justice, and fairness.
- iv. Finally, if we are going to be investing ethical and political rights/values in a National AI Research Resource, we need to include research into values/rights so they apply to novel environments. For example, what does privacy mean? What does transparency mean? What does autonomy mean?

h. *A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector;*

- i. We need to partner with industry and academia to train highly skilled AI engineers and technicians. Extensive adult education programs, distance education programs, and online courses should be developed to fit personal circumstances and schedules of those interested in pursuing AI careers. Strong programs that attract students to AI from early stages (high school and undergraduate)

and promote AI careers will be a key ingredient for growing a workforce with advanced AI expertise and sustaining the NAIRR.

2. Which capabilities and services (capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;) provided through the NAIRR should be prioritized?

- A. Open AI platforms and resources: a vast interlinked distributed collection of “AI-ready” resources (such as curated high quality datasets, software, knowledge repositories, testbeds for personal assistants and robotics environments) contributed by and available to the academic research community, industry, and government.
 - a. Provide the AI community with substantial experimental resources for both basic research and applications.
 - b. Reduce redundancy of effort and cost in the research enterprise, so research projects do not have to build up capabilities or collect new data from scratch each time.
 - c. Reduce the cost of individual research programs to integrate relevant capabilities or to compare their work with others.
 - d. Reduce cost of large teams of collaborators by providing already integrated or easy-to-integrate infrastructure.
 - e. Provide the AI community with open resources that will bolster not only research in all of academia, but also in small technology companies, companies in other sectors, and government organizations.
- B. Incentivize emerging interdisciplinary AI areas: initiatives to encourage the research community to work in interdisciplinary AI studies—e.g., AI safety engineering, as well as analysis of the impact of AI on society—will ensure a workforce and a research ecosystem that understands the full context for AI solutions.
- C. Address AI and the future of work: these challenges are at the intersection of AI with other disciplines such as economics, public policy, and education. It is important to teach students how to think through the ethical and social implications of their work.
- D. Train highly skilled AI engineers and technicians: support and build upon the National AI Infrastructure to grow the AI pipeline through community

colleges, workforce retraining programs, certificate programs, and online degrees

- a. Develop AI curricula at all levels: guidelines should be developed for curricula that encourage early and ongoing interest in and understanding of AI, beginning in K-12 and extending through graduate courses and professional programs.
- b. Create recruitment and retention programs for advanced AI degrees: including grants for talented students to obtain advanced graduate degrees, retention programs for doctoral-level researchers, and additional resources to support and enfranchise AI teaching faculty.
- c. Engage underrepresented and underprivileged groups: programs to bring the best talent into the AI research effort.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

- A. It is not just research and development that needs ethical attention, but the deployment and use are equally, if not more, important. It is imperative to incorporate ethics and related responsibility principles as central elements in the design and operation of AI systems.
 - a. Make sure that AI systems align with human values and norms to ensure that they behave ethically, taking into account potential risks, benefits, harms, and costs. In order to do this, AI systems will have to incorporate complex ethical and commonsense reasoning capabilities that are needed to reliably and flexibly exhibit ethical behavior in a wide variety of interaction and decision making situations.
- B. An AI system must be able to explain its rationale to the team members (e.g., why it suggested certain experiments) and it must make its level of uncertainty clear in a way that team members can truly understand. Given the critical applications and outcomes of modern AI, these systems must also act reliably.
- C. Explicitly codifying best practices in ethical behavior, conduct, and inclusiveness in academic, industry, and government organizations, so that inappropriate interactions, isolation, and implicit biases are eliminated from the school and the workplace.
- D. Novel computing technologies often improve our lives, but they can also affect them in ways that are harmful or unjust. It is important to teach students and practitioners how to think through the ethical and social implications of their work.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

- A. One building block that already exists for the NAIRR is the National Science Foundation (NSF) CloudBank, a cloud access entity that helps the computer science community access and use commercial clouds for research and education by delivering a set of managed services designed to simplify access to commercial clouds
- B. It is important to recognize that industry and academic work to support multicloud and hybrid cloud scenarios, since they will be enduring for some time (e.g. because of data sovereignty and privacy issues, to prevent lock-in to single-hyperscaler solutions, to take advantage of specialized cloud attributes, and to recognize the need to exploit local computing for cost and latency optimization).
- C. As discussed in a recent [CRA-Industry roundtable](#) it is important to recognize that commercial clouds provide capabilities that could not be rivalled by government-funded alternatives. To avoid wasteful investment, alternative strategies need to be considered.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

- A. The NSF CloudBank suggests one direction for public private partnerships, which is mainly to get the hyperscalers to actively contribute or subsidize lower-cost services for government and academic research. And per above, there is another kind of public-private partnership that would help drive initiatives in multi-cloud/hybrid cloud access that would help prevent single-cloud lock-in and help perpetuate dual-source options for government/academic procurement.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

- A. The US education system currently makes use of computing technologies, some of which are enhanced by AI, but there is still great room for improvement of these technologies, and even greater opportunity for full adoption to enhance our education system if access were free for all on all technical levels. If students had free access to high quality education, lifelong education and training would only

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Consumer Reports

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1, 2021

National AI Research Resource Task Force
Attn: Ms. Wendy Wigen, NCO, NITRD Program
2415 Eisenhower Avenue
Alexandria, Virginia 22314

Re: Request for Information and Comment on a National Artificial Intelligence Research Resource

Dear Members of the National AI Research Resource Task Force:

Consumer Reports (CR) writes today in response to the Request for Information and Comment on a National Artificial Intelligence Research Resource. Consumer Reports is an expert, independent, non-profit organization whose mission is to work for a fair, just, and safe marketplace with and for all consumers and to empower consumers to protect themselves.¹ We applaud the The Office of Science and Technology Policy and the National Science Foundation creating a shared research infrastructure that would provide artificial intelligence (AI) researchers and students across scientific disciplines with access to computational resources, high-quality data, educational tools, and user support. (Question 2) Smaller companies, academics, and public-interest researchers do not always have the resources to develop larger and more complicated AI models — NAIRR should prioritize providing things like cloud storage and computing capacity to these groups. Democratizing AI can not only lead to more fair outcomes for affected populations but also can mitigate harm done by biased or otherwise detrimental algorithms.

Technology that uses AI has the potential to discriminate across a wide variety of sectors and applications. Our concerns about the use of AI are not unique to technology. They are about fairness. AI, when training data is biased, or when algorithms are flawed due to human biases, can reproduce and further entrench existing harms, or create new ones. As AI becomes more integrated

¹ CR works for pro-consumer policies in the areas of financial services and marketplace practices, antitrust and competition policy, privacy and data security, food and product safety, telecommunications and technology, travel, and other consumer issues in Washington, DC, in the states, and in the marketplace. Consumer Reports is the world's largest independent product-testing organization, using its dozens of labs, auto test center, and survey research department to rate thousands of products and services annually. Founded in 1936, Consumer Reports has over 6 million members and publishes its magazine, website, and other publications.

into everyday products and daily life, it is important that its development be democratized and accessible to all in order to mitigate harmful effects.

Discrimination in algorithms is a serious concern and has the potential to erode much of the progress made by U.S. civil rights law. There are many sources of bias in algorithms, but a significant way algorithms produce discriminatory outputs is due to biases that stem from societal inequities. For example, Black communities tend to be overpoliced so a disproportionate percentage of crime data is collected from these communities; when an algorithm is designed to predict where crimes occur more often in a particular city in order to better allocate policing resources, for instance, it could point to the Black communities that are already being heavily policed.² There are many other sources of biases in algorithms during the design process including other biased data collection methods, the specific type of model being used, as well as the attributes of the data the engineer chooses as being important to the final outcome.

(Question 3) It is important that inclusive datasets that more fully represent the populations the algorithm is trying to make predictions or classifications for are available to the public. Often, private companies, particularly smaller ones, do not have the resources to perform proper data collection and must resort to open-source databases that tend to be of lower quality or incomplete. Also, public-interest researchers attempting to audit or reverse engineer potentially harmful algorithms are not able to do so without higher-quality training data.

(Question 5) NAIRR can mitigate this issue by partnering with private companies who have more complete datasets to provide data to the public, or sourcing data from different locations and testing it to ensure completeness and accuracy before making it publicly available. Furthermore, NAIRR should be testing the data across different dimensions like protected classes like race, gender, etc. to ensure these demographics are adequately represented and provide markers to their datasets when they are not. They should also release guidelines for how to test for completeness in training data and for what applications this would be particularly useful or necessary. Finally, they should also allow for the public to contribute to these datasets; treating datasets as an open-sourced tool can democratize the AI development process and encourage competition when smaller companies or individuals are able to get access to robust and high-quality data.

(Question 3) AI educational tools are necessary when developing fair and inclusive technology. Responsible research and ethics are not always at the forefront of early-stage companies, and providing resources that can help companies think through complex social issues is vital when mitigating AI harm and maximizing its benefits. NAIRR should also perform research on and release guidelines regarding the potential misuse or misapplication of AI. For example, the use of pseudoscience and physiognomy are on the rise in AI applications; some companies claim that

² O'Donnell, Renata M. "Challenging Racist Predictive Policing Algorithms Under the Equal Protection Clause," *NYU Law Review*, 2019, <https://www.nyulawreview.org/wp-content/uploads/2019/06/NYULawReview-94-3-ODonnell.pdf>.

their AI can do things that are not necessarily possible or substantiated by science.³ NAIRR should make clear why certain uses of AI are harmful and/or misleading to discourage companies from creating these sorts of models. This research should be done in conjunction with social scientists, AI researchers and ethicists, and civil rights groups. Research should also focus on privacy-protecting methods that allow for careful examination of how civil rights can be potentially impacted by AI without disclosing anyone's personal information.

(Question 1) Transparency is an important tool that NAIRR should be encouraging builders of AI technology to leverage in order to mitigate harm. NAIRR should perform research on algorithmic impact assessments and provide guidance on how companies should be testing for bias and reporting it to appropriate parties. This includes disclosure of data used in the algorithm, an explanation of how the algorithm works, the steps the company took to test for disparate impacts, and how they mitigated harmful effects if identified. NAIRR should also perform research on auditing techniques and release guidelines for potential third-party auditing. This includes information like what sorts of algorithms should be subject to an audit, how that audit should be carried out and what entities can perform it, and what kinds of information companies should give to auditors to perform a successful audit; this may entail working with other agencies and/or private auditing groups to provide some sort of accreditation process for audits. Transparency should also be integral to NAIRR itself. All research done by NAIRR and all partnerships and stakeholders for any NAIRR project should be disclosed to the public. NAIRR should primarily be focused on researching AI that is beneficial to the public and strategies to mitigate or avoid harm.

We are excited about this new initiative and thank OSTP and the NSF for creating this task force. While AI has the potential to do good, its potential harms are severe and can infringe on Americans' civil rights. Our suggestions can help ensure that AI research and development becomes more democratized which will mitigate harm caused by this emerging technology. Thank you for your consideration.

Sincerely,
Nandita Sampath
Policy Analyst

³ Narayanan, Arvind. "How to Recognize AI Snake Oil."
<https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

CrowdAI

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



RFI Response: National AI Research Resource (September 2021)

- 1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]**

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

We recommend incorporating end-to-end (i.e., from data ingestion and labeling to model production), open, modular, code-free AI enablement tools in order to facilitate access to the proposed NAIRR beyond just AI/ML engineers (of which there is short supply).

Such tools provide numerous benefits at the enterprise-level, including but not limited to:

- Scale the adoption of AI by enabling a wider workforce to engage in AI research and AI-enabled research without needing to turn vast numbers of people into AI/ML engineers or data scientists
- Help standardize data ontologies and curate data sets for wider, more efficient use
- Provide a single user-interface
- Streamline secure access controls

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

AI models require case-specific training data, and cannot be tuned without



RFI Response: National AI Research Resource (September 2021)

representative data that is appropriately labeled. However, some USG programs expect to evaluate and purchase models without providing AI vendors consistent, relevant, and diverse training data.

AI model accuracy directly results from robust, well-curated training data. Without consistent, relevant, and diverse training datasets, AI vendors are more likely to over-fit or under-fit to the data, decreasing the quality of models.

Part of the problem has to do with the USG operating under a misperception that if it provides all available training data to AI vendors, they will be "giving too much away," and the USG will lose the ability to effectively test and evaluate delivered models. While withholding some data to support testing and evaluation generally is considered a best practice for AI development, we often have experienced the USG not providing adequate training data to build truly scalable, domain-specific models.

Another part of the problem stems from the security classifications of data.

We recommended the following solutions to these barriers to the dissemination and use of high-quality government data sets:

- NAIRR should commit to providing AI researchers without security clearances access to unclassified training data.
- NAIRR should examine the feasibility of holding clearances for cleared researchers who are affiliated with uncleared organizations and would like to work with classified data.
- NAIRR should coordinate a USG effort to put forth a standardized ontology for creating training data.

1. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?



RFI Response: National AI Research Resource (September 2021)

We recommend incorporating end-to-end (i.e., from data ingestion and labeling to model production), open, modular, code-free AI enablement tools in order to facilitate access to the proposed NAIRR beyond just AI/ML engineers (of which there is short supply).

Such tools provide numerous benefits at the enterprise-level, including but not limited to:

- Scale the adoption of AI by enabling a wider workforce to engage in AI research and AI-enabled research without needing to turn vast numbers of people into AI/ML engineers or data scientists
- Help standardize data ontologies and curate data sets for wider, more efficient use
- Provide a single user-interface
- Streamline secure access controls

2. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

We realized early on the importance of building AI that responsibly captures our users' workflow throughout the entire AI development process. Keeping the researchers' workflow in mind is the best way to build AI that provides the most value, and reinforces principles of ethical and responsible AI R&D.

Tools that engage users throughout the process from the beginning to the end, and back again (that is, for model sustainment), serve as guardrails for making sure people are building AI models in a manner that is responsible, equitable, traceable, reliable, and governable.

Capturing researchers' workflow can best be done through end-to-end (i.e., from data ingestion and labeling to model production), modular, code-free AI enablement tools.



RFI Response: National AI Research Resource (September 2021)

There are multiple pieces to the AI process that need to be taken into account when building a responsible AI ecosystem. These include considerations around -- and accompanying enablement tools for --

- Compute
- User Management
- Data (ingestion, management, labeling)
- Model Training
- Model Testing & Evaluation
- Model Deployment & Sustainment

Below are considerations around each of these pieces of the AI process:

- Compute - AI, especially computer vision (or CV), requires significant compute resources.
 - Deciding where to allocate compute resources should be considered from a mission perspective. For example, within the GEOINT world, identifying Indications & Warnings, or doing mapping wildfire extent, requires real-time analysis. To reduce latency of model outputs, the reliable option here from a mission perspective would be to allocate compute resources as close to the sensor(s) as possible. For other missions, real-time analysis might not be as pressing, so compute resources can be allocated closer to end users.
 - That said, cost also factors here when considering scalability.
- User Management - While software is integral to AI, at the end of the day, it's people that are interacting with data.
 - It's important to have a user management tool to grant and restrict access to 1) see and/or edit data, and 2) train, test, and deploy models, in a manner that accords with data protection.
- Data - Data forms the backbone of AI. It's what it's built on, and it's what makes an AI model performant, useful, and responsible, or not.



RFI Response: National AI Research Resource (September 2021)

- Researchers need to know why they need plentiful, diverse, and appropriately labeled data in order to build robust models, which requires easy-to-use (i.e., code-free) tools to ingest data and label it.
- Data standards for ontology and metadata would go a long way when it comes to an AI ecosystem's governability and reliability. NAIRR and its AI stakeholders should work together to set standards for ontology and metadata, would result in the kind of consistency and quality we would like to see and would benefit everyone. We've seen this being useful for FEMA disaster response categories.
- Data management tools can help mitigate bias as well. E.g., by grouping data with similar data with respect to any given characteristic (e.g., data from a certain type of sensor), and then flagging for the user than if all the training data for a model comes from only one grouping (e.g., all of the data comes from just one sensor) it can't be expected to generalize, or perform on data beyond that type (in this case data from other sensors).
- Model Training - Model training needs to be transparent and traceable when it comes to how a model was trained, and on what data. to have the ability to log model training, especially if AI is to earn the long-term trust of users.
- Model Testing & Evaluation - Here it's important to have visual tools and education around test & eval issues to ensure models are reliable.
 - E.g., researchers need to understand that test set data must be separate from training data.
 - Researchers need to understand how their particular use-case informs the precision and recall they might be looking for (e.g., in the case of countering illicit trafficking there may be a tolerance for lots of false positives, but absolutely no false negatives)

RFI Response: National AI Research Resource (September 2021)

- Model Deployment & Sustainment - In order for models to be reliable for mission use, there needs to be an easy way for users to deploy them and continuously sustain them through retraining with new data and new contexts.
 - Here, a tool to automate deployment through containerization is key.
 - For sustainment, a feedback loop wherein users rate model output, thereby creating more training data for the model to improve with.

In addition, intuitive training sessions, workshops and/or documentation on data literacy are extremely important. Researchers need to know why they need plentiful, diverse, and appropriately-labeled data in order to mitigate bias.

We recommend conducting workshops for users, in addition to the brief tutorials that may exist in the code-free enablement tools for data ingestion, data management, and data labeling.

A data literacy workshop or tutorial would typically highlight why data needs to not only be plentiful, but diverse as well in order to create a model that is better generalized to the task at hand, and not overfit on a small homogeneous set of data where the model might pick up on and be biased towards the wrong characteristics.

A simple example we might use to highlight this: a model to detect individuals who are armed with weapons is only trained on dark-skinned individuals with weapons and light-skinned individuals from different video feeds without weapons, the model might erroneously assume that the defining characteristic of having a weapon is having dark skin.

Another simple example: show two sets of pictures, 6 cats and 6 dogs. In the pictures, it is clear to humans that there are 6 cats and 6 dogs. The six cats are on sofas and the six dogs are on the grass, the animals take up less than 10% of the image. Then show a dog on a couch and a cat on grass. The model may get confused, because models take the easiest approach.



RFI Response: National AI Research Resource (September 2021)

The data literacy workshop or tutorial would also cover how to label data appropriately for the task at hand.

Take for example an object detection bounding box model. It's important to show what's an appropriately sized bounding box to use in any given piece of data to turn that into good training data (the bounding box needs to capture all of the relevant pixels and nothing else).

3. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Given 1) the relative scarcity of AI and machine learning expertise in the workforce, 2) the need for solutions that recognize the iterative nature of effective AI, and 3) the need to create and manage quality training data for AI models, we're seeing the emergence of no-code enablement tools and platforms. Lots of companies are offering components of the AI pipeline, like data labeling or data management tools.

At CrowdAI, we focus on offering every component of the end-to-end AI pipeline, specifically as it pertains to computer vision (user management, data labeling/management, model training, model testing and evaluation, and deployment)

4. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Companies with end-to-end, modular, code-free AI enablement tools should offer use of their platforms at discount or no-cost, similar to academic partnerships they might already have.



RFI Response: National AI Research Resource (September 2021)

5. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

There is a short supply of AI skills within the USG, and private sector. Obtaining the required skills to develop AI solutions is onerous. The skills gap often hinders the USG's ability to purchase, deploy, and/or operate AI solutions.

The only way to efficiently democratize access to AI R&D is to arm the workforce with code-free AI enablement tools through which they can build, deploy, and sustain their own AI models, thereby removing reliance on the short supply of AI engineers.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

**Jared Freeman, Drew Leins, Niall
Gaffney**

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

National AI Research Resource An RFI Response

Jared Freeman, Drew Leins, Aptima, Inc.
Niall Gaffney, Texas Advanced Computing Center

Corresponding author: [REDACTED]

1 September 2021

[Introduction](#)

[Response to RFI Questions](#)

[1. Roadmap Options](#)

[Goals](#)

[Ownership & Administration](#)

[Governance & Oversight](#)

[Capabilities](#)

[Data Dissemination](#)

[Security](#)

[Privacy & Civil Rights](#)

[Sustenance](#)

[Organizational Structure & Process](#)

[2. Capability Prioritization](#)

[3. Ethical AI](#)

[4. Building Blocks](#)

[5. Public-Private Partnerships](#)

[6. Democratization](#)

Introduction

This response to [the NAIRR RFI](#) is inspired by the authors' experiences planning and/or executing large-scale, distributed research programs in which AI and data for AI are the primary technical products. These programs are DARPA's Synergistic Discovery and Design (SD2), which aspires to accelerate the pace of exploration and production of biological and chemical constructs through AI, and NIH's Bridge to Artificial Intelligence (Bridge2AI), which is designed to develop flagship data sets

that advance AI applications to biomedical problems. The authors of this response are primes for evaluation (Aptima) and computing infrastructure (TACC) on SD2, and co-proposers on Bridge2AI.

Response to RFI Questions

1. Roadmap Options



This section responds to question #1: What options should the Task Force consider for any of roadmap elements A through I above, and why?

A. Goals

This section concerns: A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

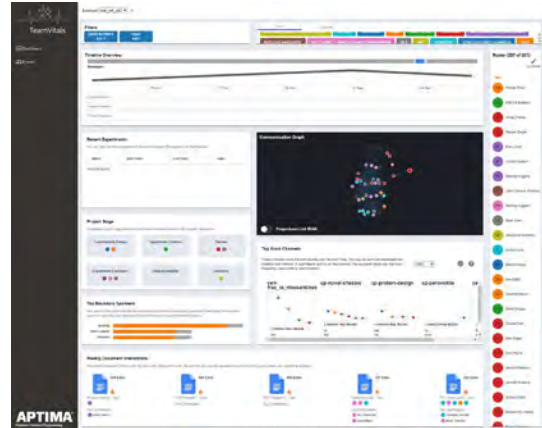
The NAIRR should set as a goal transparently and measurably advancing science by creating technologies and data that are efficient, usable, used, and valued by an engaged community of researchers and application specialists. To achieve this goal, NAIRR should apply metrics and measurement technologies that quantify its progress. Specifically, NAIRR should:

- **Apply metrics of scientific progress** -- Define classes of metrics that are relatively independent of scientific domain. Define measures that are simple to comprehend, such as multiplicative improvement over baseline. Solicit measurements both of baseline and of effects. In DARPA SD2, this strategy revealed striking effects of the program's sociotechnical system on speed, quantity, and accuracy of research or research constructs, increased complexity of constructs (e.g., synthetic circuits in biological systems), and reductions in human level of effort.
- **Apply metrics of system efficiency** -- Instrument systems that extract, transform, and load (ETL) data. Identify long latencies and cyclical processing, diagnose the root causes, and revise these data and their metadata to improve quality or revise the systems to make them more robust to poorer quality data. In DARPA SD2, this strategy helped the program to improve its systems, its data products, and the pace of discovery and design.
- **Apply metrics of use and utility** -- Instrument NAIRR systems to measure use (e.g., downloads, pulls, pushes) of the products it offers. Elicit feedback (e.g., through user surveys) concerning



other aspects of those products. In DARPA SD2, such data were used to profile technical products (data, software) with respect to comprehensibility (documentation), accessibility (e.g., public posting), utility/value, and usability (e.g., support required to modify the product or put it to use). The traditional TRL levels fail to capture these attributes.

- **Apply metrics of organization & community evolution** -- Instrument NAIRR collaboration platforms and conduct periodic surveys to capture data concerning its composition (e.g., professional and cultural diversity of personnel), organizational structure, process, member engagement, and climate. Perform computational and statistical modeling to identify positive and negative trends relative to best practices of organizational science. Identify correlational and causal effects of these factors on S&T productivity. Reward laudatory trends. Intervene to address worrisome trends, relative to the organization charter. Identify and leverage emergent, informal “boundary spanners” (individuals who are central to multiple groups, and emergent communities whose work has high value (figure). DARPA SD2, the USAF, and Army have applied these techniques in training and operations, and have funded and government-owned technologies to automate them.



B. Ownership & Administration

This section concerns: B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including: i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

- **Ensure equity, diversity, and inclusion** -- Administration of NAIRR should commit in its charter, realize in its membership, and codify in its processes a commitment to Equity, Diversity, and Inclusion. Measures of these attributes should be established (see Governance, above) to assess the baseline, current state, and evolution of the organization with respect to EDI. For example, inclusion can be assessed from records of participation in meetings, chat, and collaborative documents. Diversity can be assessed by comparing the distribution of members over target attributes (e.g., race, socio-economic status, profession) to norms (e.g., the distribution of the US population, of the scientific population). Equity can be assessed through analysis of the distribution of power, measured as roles, control of resources, etc.

C. Governance & Oversight

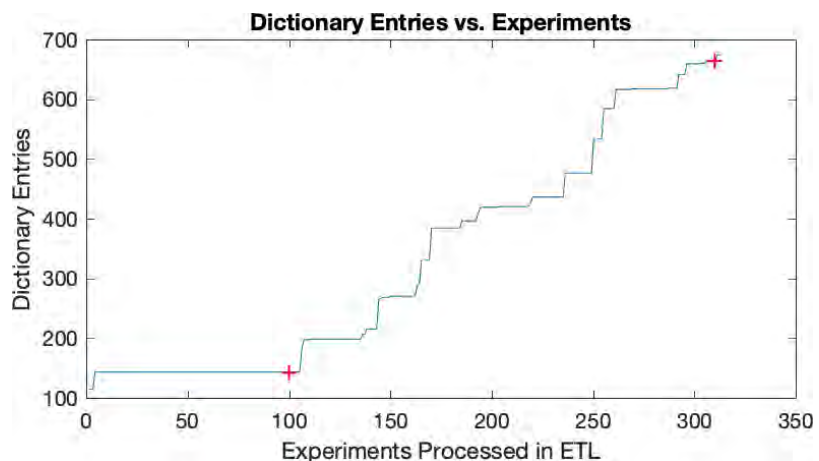
This section concerns: C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

- **Design governance for productivity** -- A Steering Committee should be established to set the goals of the NAIRR and assess its progress. Its members should include the administering organization, leads for major groups within NAIRR, and key representatives of the Accountability Council (below).
- **Design governance for accountability** -- An Accountability Council should be established whose members are sampled from the population of data providers (e.g., human research participants, the National Archives for historic records, bio-medical materials developers), data processors (e.g., data scientists, developers of ETL and analytic tools), consumers of derivative data (e.g., physicians, physicists, economists), and ethicists. This panel should identify issues in the conceptualization (e.g., intended use, sampling), generation, storage, processing, and use of AI data and tools; set policy regarding these; and review potential violations of that policy. As noted by Dr. Francis Collins, Director of NIH, in testimony to Congress in 2021, inclusive governance raises public trust and enhances recruitment and retention to research efforts.

D. Capabilities

This section concerns: D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

- **Invest in metadata design** -- Support collaborative design of metadata that describe data and the context in which they were created, such as the experimental



participants and/or materials, procedures, analyses, and relationship to published findings. In DARPA SD2, such metadata stores were correlated with, and likely causal, to dramatic increases in the quantity of biological research (analysis by Jacob Beal, BBN, for SD2; see figure).

- **Invest in multi-entry workflows** -- Support development not just of independent tools but of workflows, typically enabled by a framework for integrating tools in novel combinations that are fit for a specific purpose. Equally important is that such workflows enable users to access data at multiple points in processing and storage, to support users with different levels of technical skill and different research objectives.

E. Data Dissemination

This section concerns: E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

- **Address regulatory restrictions** -- HIPPA, Institutional Review Boards, tribal law, and other regulations restrict the distribution of raw data concerning human subjects. Datasets should be clearly labeled with such restrictions, and restrictive documents (such as informed consent forms) should be made available. Access to restricted data should be conditioned on NAIRR review (e.g., that a local IRB has approved processing of human subjects data; that a researcher has tribal authorization to analyze data concerning an American Indian tribe). Mechanisms should be available through NAIRR to anonymize human subjects data, and agreements required that users will not reverse engineer human identification.
- **Document data** -- Data are always collected with intent, process, and scope. These are key elements to achieving leveraging any dataset to produce results. While there are many successful AI outcomes driven by data repurposing (e.g. twitter flu predictions) for any AI data or model to be accurately leveraged, these key elements of metadata must be well described and revealed. Data and models should be published with documentation of their intent, process, and scope. In a perfect world, one would also capture both successful insight as well as the failures in using any data or combining any data sources to both promote accurate outcomes from data and to help inform future data generation efforts of limitations their intent, process, and scope may have on its utilization in addressing future questions they or others may want the data to address.
- **Automate data documentation** -- Data dissemination relies on both data discovery, tagging, interpretation, and linkages to other data. This requires effective metadata behind the data. Currently most systems require well defined schemas, ontologies or constrained vocabularies which creates a significant burden on the human data producers, scientists, and archivists who must keep these “standards” evolving as the utility and linkages of the data evolve with changes and discoveries. Lack of support often leads to degraded utility of that data. Recent works have shown some well managed NLP/AI systems can create such linkages through published literature (c.f., DIVE by Xu, Gupta, Jaiswal, Taylor, and Lockhart, 2016). However, this has not been extended to the multitude of data sources available. Support must be given to automate the evolution of metadata systems to keep key data relevant and useful for improving outcomes from federally funded projects and programs.

F. Security

This section concerns: F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;

- (No response).

G. Privacy & Civil Rights

This section concerns: G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

- (No response).

H. Sustenance

This section concerns: H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector;

- (No response).

I. Organizational Structure & Process

This section concerns: I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

- (No response).

2. Capability Prioritization

This section responds to question #2: Which capabilities and services provided through the NAIRR should be prioritized?

- (No response).

3. Ethical AI

This section responds to question #3: How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

- (No response).

4. Building Blocks

This section responds to question #4: What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

- **Learn from model programs** -- NAIRR should study extant programs for models of governance, capabilities, and productivity. Among the candidates are:
 - DARPA SD2 (PM: Joshua Elliott) -- The Synergistic Discovery and Design (SD2) program developed data-driven methods to accelerate scientific discovery and robust design in domains that lack complete models.
 - NIH Bridge2AI -- The NIH Bridge to Artificial Intelligence (Bridge2AI) Program seeks to bridge the biomedical and behavioral research communities with the rapidly growing community of experts developing AI/ML models by producing flagship datasets that adhere to the FAIR principles (Findable, Accessible, Interoperable, Reproducible) and critically integrate ethical considerations in preparing data for computation.

5. Public-Private Partnerships

This section responds to question #5: What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

- (No response).

6. Democratization

This section responds to question #6: Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

- (No response).

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

The Data Foundation

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1, 2021

Wendy Wigen
2415 Eisenhower Avenue
Alexandria, VA 22314, USA

Re: Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource:

Ms. Wigen:

Thank you for the opportunity to provide comments on the implementation plan for a National Artificial Intelligence Research Resource.

The Data Foundation is a non-profit organization that seeks to improve government and society by using data to inform public policymaking. Our Data Coalition Initiative operates as America’s premier voice on data policy, advocating for responsible policies to make government data high-quality, accessible, and usable. The Data Coalition supports policies that encourage responsible, ethical deployment of innovative and emerging technologies that foster rigorous data analysis for improved decision-making, while facilitating equitable data use and appropriate privacy protections.

Artificial intelligence and machine learning technologies can have real, tangible benefits for the research community, but that requires access to large amounts of high quality data. AI researchers may need data from multiple government agencies, different levels and branches of government, as well as non-profit or private sector organizations. Our current data infrastructure is decentralized and fragmented, and large portions of it are out of date, unable to keep up with the technical demands of privacy and data sharing that could be beneficial for research on AI and research conducted with AI.

With that in mind, the Data Coalition recommends the task force explore how to create a modern data access and linkage infrastructure for researchers inside and out of the government, in particular the creation of a National Secure Data Service. Both the 2016 Commission on Evidence Based Policymaking¹ (Evidence Commission) and the Committee on National Statistics² recognized the need to improve data access and strengthen privacy protections and recommended establishing a new entity in the federal government to support data linkage and combination activities, mechanisms for

¹U .S. Commission on Evidence-Based Policymaking (CEP). The Promise of Evidence-Based Policymaking: Final Report of the Commission on Evidence-Based Policymaking. Washington, D.C.: Government Publishing Office, 2017a.

² National Academies of Sciences, Engineering, and Medicine (NASEM) Committee on National Statistics (CNSTAT). Innovations in Federal Statistics: Combining Data Sources While Protecting Privacy. Washington, D.C.: The National Academies Press, 2017a. Available at: <https://doi.org/10.17226/24652>.

improving data access, and enhancements for privacy with new and emerging technologies, all of which would benefit AI research.

Several proposals have been circulated on ways in which a National Secure Data Service can be established, consistent with the recommendations of the Evidence Commission. Four such proposals are outlined in a [2020 report from the Data Foundation](#) on design considerations for such a service, including legal authority for privacy protections, ability to access and acquire data, as well as scalability and sustainability.³ Additional efforts to establish a data service are underway in Congress, with the introduction of the National Secure Data Service Act ([H.R. 3133](#)). The Advisory Committee for Data for Evidence-Building is also working to make recommendations on how to facilitate data sharing, data linkage and privacy enhancing techniques, with publication expected later in the year.

We encourage the task force to consider how to sync efforts to create an Artificial Intelligence Research Resource with existing efforts to create a National Secure Data Service as a way to address the structural challenges to improving access to data in a secure setting.

Thank you for the opportunity to provide feedback on plans for an Artificial Intelligence Resource. We welcome any questions on our comments and look forward to supporting your efforts in this area.

Sincerely,

Corinna Turbes
Policy Director

³ Hart, N. R., & Potok, N. (2020). (rep.). *Modernizing U.S. Data Infrastructure: Design Considerations for Implementing a National Secure Data Service to Improve Statistics and Evidence Building*. Washington, DC: Data Foundation. <https://www.datafoundation.org/modernizing-us-data-infrastructure-2020>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Deloitte

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Implementing a National Artificial Intelligence Research Resource

October 1, 2021

In response to the Implementation Plan for a National Artificial Intelligence Research Resource Request for Information (RFI)

RE: RFI Response: National AI Research Resource

Deloitte Consulting LLP (Deloitte) is pleased to submit to the Office of Science and Technology Policy (OSTP) and National Science Foundation's (NSF) Request for Information to implement the National Artificial Intelligence Research Resource (NAIRR) Implementation Roadmap.

Deloitte believes the NAIRR initiative fills a critical and strategic gap in the government's AI initiative. To support its success and address all RFI parameters, our response draws on Deloitte's AI expertise in both the government and private sector domain. This includes our in-house data science talent, experts across AI and government, experience developing AI solutions for public and private sector clients, and familiarity with the perspectives of AI researchers through regular surveys and our professional services.

We remain ready and interested to support the NAIRR Task Force (TF). We would be pleased to share our experience advising and implementing AI across government and commercial with OSTP and NSF. Should you have any questions, please contact me at ([REDACTED]).

Sincerely,

Ed Van Buren
Principal
AI in Government Leader
Deloitte Consulting LLP

Table of Contents

Company Profile

Company Name	Deloitte Consulting LLP
Headquarters Location	New York City
Contact Name	Ed Van Buren
Contact Title	Principal, Deloitte Consulting LLP
Contact Email Address	[REDACTED]
Contact Phone Number	[REDACTED]
Primary Type of Service(s) Provided	Software Development, Professional Services, Management Consulting, Technical Support, Maintenance, and Support Services

Executive Summary

Artificial intelligence holds great promise for U.S. economic growth and prosperity as well as national competitiveness, more generally. However, the evolution of AI has been inhibited by issues of data access, security, and quality; the increasing computational demands and complexity of AI modeling; and differing interests and disparate resources of various AI constituencies in the private sector, government, and academia.

These issues have routinely emerged across industries. In our annual [State of AI in the Enterprise](#) survey, we observed varying rates of AI adoption, varying preferences in AI development approaches, and varying competencies for effective AI procurement between private sector and public sector respondents. We recently developed an [AI Dossier](#) to help leaders understand when, where, and how to deploy AI within their organization. The Deloitte AI Institute for Government collaborates with an ecosystem of academic institutions and thought leaders to identify practical ways to develop and deploy AI.

While our response details considerations and pathways for the establishment and sustainment of the NAIRR, we highlight three areas the TF will need to navigate: 1) need for openness & transparency to foster cutting-edge innovation while adhering to security demands of sensitive

and proprietary data; 2) need to collaboratively address challenges of interoperability across data types and models; and 3) need to foster rapid innovation and R&D with a focus on ethical and trustworthy AI as well as diversification of AI developer talent.

As detailed below, we believe the solution lies within the NAIRR governance model, its resource base, and stakeholder inclusion so AI research is not misunderstood and is widely beneficial.

Responses to RFI Questions

Q1: What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

	Response
A.	<p>The TF should prioritize a flexible development process, and employ an agile methodology, allowing lessons learned to be incorporated iteratively. Defining success criteria from the onset will be critical to a successful implementation. The TF should consider Key Performance Indicators (KPIs) based on goals of the NAIRR and develop yearly metrics targets, even if preliminary and directional, to drive and evaluate success. Illustrative metrics could include: number of projects executed, number and variety of quality dataset uploads and downloads, activity levels on knowledge portals, computational resource availability and sharing, new users onboarded, and partnerships established. Additionally, the NAIRR can capture demographic data to assess if underserved communities, institutions, and regions are accessing and utilizing the NAIRR.</p>
B. & C.	<p>The NAIRR’s success will depend on the organization responsible for program implementation, deployment, and management. Considerations for the TF during the selection process could include:</p> <ul style="list-style-type: none"> • Mission Alignment: The goal of “democratizing access to the cyberinfrastructure that fuels AI research and development” should fit the organization’s mission - this facilitates shared purpose and impact. • Positioning: The organization should have industry agnostic scope to allow research across industry, government, and academia. An agnostic approach will help with and contribute to an equitable partner experience. • Authority: The organization should have resources and capacity to establish direction, make programmatic decisions, and manage resources. <p>To maximize adoption and utilize existing resources and expertise, the NAIRR should consider a hub and spoke governance model. This would centralize the NAIRR as a hub that governs the standup and drives policy, standards, and driving adoption, while coordinating with existing agency missions, partnerships, engagements, infrastructures, and research communities (spokes). As the central hub, the NAIRR could establish a PMO charged with strategic direction, programmatic decision making, and resource allocation. This approach allows each research spoke to be established based on individual requirements and provides them autonomy which can lead to innovation. The NSF AI</p>

	Institutes Program is a good example of hub and spoke governance where the NSF is a central hub and the partnered universities are the spokes.
D.	<p>To democratize AI R&D capabilities, the TF will need to prioritize foundational capabilities available in the initial operating capability. This includes:</p> <ul style="list-style-type: none"> • Data infrastructure with pretrained models and quality datasets • Hybrid-cloud compute infrastructure, compute power and services • Knowledge management portal for educational tools and services • Proper testing & evaluation capabilities to allow models to be validated • Change Management and strategic comms to drive engagement and adoption <p>Please see 3.2 for a full response.</p>
E.	<p>The ability to share and disseminate high quality data across the NAIRR is contingent on the successful mitigation of several operational and technical concerns. Each will require discrete risk mitigation techniques to overcome the concerns.</p> <ul style="list-style-type: none"> • Security Concerns – The NAIRR should create a framework to detect security and cyber risks and develop incident response plans in the event of a compromise. • Usage Concerns – To avoid data confusion / misinterpretation, each dataset should come with an abstract or documentation to describe the data and its formatting. • Attribution Concerns – License requirements for data uploaders to populate can be combined with an auto-attribution feature from the uploader’s profile to ensure that if institutions reuse datasets / resources, they cannot remove the license or text that is copyrighted. • Cultural Barriers – Siloed mindsets can inhibit progress in research, data, and models. To encourage data sharing, institutions can be rewarded and credited with additional resources (i.e., storage, compute, etc.) for their contributions.
F.	<p>To successfully operate and house research from various research hubs, the program must invest in state-of-the-art technologies to encrypt data and manage access to its resources, to include data, models, and related research. To identify appropriate security and access management controls, the NAIRR could work with National Institute of Standards and Technology (NIST) to curate standards to address security and operational constraints (i.e., NIST 800-53). This will help the NAIRR develop a resilient federated Identity and Access Management (IdAM) control system across its hybrid-cloud infrastructure.</p> <p>To enable this, the NAIRR should consider a zero-trust architecture, a data-centric security framework which requires strict identity verification for devices and personnel. An approach of least privilege can enhance the NAIRR’s cyber risk posture by limiting impact of cyber breaches and improving containment capabilities. This approach will require the orchestration of multiple organizational components, including buy-in from oversight organizations, capability providers, and research hubs.</p>
G.	<p>Trust between users and developers is key in the discipline of AI. Trust is earned over time and lost in an instant. In response to Executive Order 13859, NIST developed a framework to establish standards to bolster public trust and confidence in AI applications. These standards address societal and ethical issues, governance, and privacy.</p>

	To assess privacy and civil rights requirements associated with the NAIRR conducted research, the NAIRR should consider a framework to assist researchers with assessing elements underpinning this trust, including data attributable to consumers, citizens, or patients. The NAIRR should consider NIST’s Trustworthy AI Framework .
H.	Sustained and predictable federal funding will be essential for the NAIRR’s initiation. This will allow objectivity and independence necessary to seed / support cutting-edge innovations and access to educational resources and data. Alternative arrangements, such as subscription-based or fee-for-services may undermine the NAIRR’s ability to develop an active user base, promote innovation, or skew / narrow direction of R&D priorities. Private sector partnerships remain a potentially impactful multiplier for the NAIRR and must navigate unique interests. Prospective private sector partners will be incentivized to engage in NAIRR initiatives if they can showcase / develop internal talent, permit cooperative problem solving with public sector researchers, highlight their tools and services, or provide compute resources and technology to funded R&D activities, potentially at discounted rates if it provides them visibility and branding opportunities.
I.	<p>The NAIRR should be established in phases, with iterations to assess, evaluate, and adjust. Stakeholders should agree on capabilities needed for success. Parameters for establishment and sustainment can be broken into three steps 1) determine public authority (legal framework and political realities) 2) define project needs & objectives (i.e., speed, efficiency, degree of certainty) and 3) determine owner for each project (i.e., capabilities, financial, operational, and risk transfer).</p> <p>To help the NAIRR sustain its operations, the TF can consider an executive steering group, similar to the one operated by the Joint AI Center (JAIC). While the JAIC’s original roles and responsibilities were derived from the National Defense Authorization Act, the Agency established an executive steering group to enable stakeholders to interact with one another and align priorities. This group is composed of senior leaders and General officers, who work across areas (i.e., acquisition, workforce development, standards, AI ethics, and AI policy) to ensure conversations and perspectives are integrated across the enterprise.</p>

Q2: Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

One of the key capabilities and services the NAIRR can provide to users is **high quality clean datasets and pretrained models**. Providing resources such as clean datasets and pre-trained models allow users to quickly spin up models which can be iterated on. Availability and access to data is one of the greatest obstacles. The NAIRR should prioritize user-friendly methods to help researchers and institutions identify datasets best suited for their problems. This involves data categorization (supervised, unsupervised, cleaned, uncleaned, etc.) to start, and can be advanced into a recommender system to improve information sharing and reduce friction for finding available datasets. One example is NIH’s FAIRshake program which makes data findable, accessible, interoperable, and reusable. Furthermore, a repository of trained models shared between institutions allows for a model economy that can scale by building on each other.

With the development of deep learning models and others requiring increasing compute, the need for a **shared compute infrastructure** should be prioritized. Shared infrastructure with robust compute power is critical to the progression of AI. The easiest solution to provide compute power is established cloud and hardware vendors. To begin, the NAIRR should instantiate a single CSP, with the goal of building a hybrid on-premise and multi-cloud environment, which allows for greater interoperability. This approach reduces initial costs and programmatic complexity and allows for at-speed scalability. There are instances which could require localized processing and high-performance computers. This on-premise capability could be used for certain use cases, datasets and / or model development that would benefit from high performance compute.

The NAIRR should invest in a **knowledge management portal** to provide educational tools for new and seasoned data scientists. This portal can foster collaboration among institutes and researchers and create an open knowledge network to aid in idea creation and problem solving.

To ensure research conducted through the NAIRR's resources does not introduce bias and risk, **testing capabilities and procedures** must be in place for users when evaluating and validating their models. These tools should be flexible and allow for continuous testing, integration, and deployment, which will instill justified confidence in models produced and used.

For the NAIRR to sustain its operations, the TF could implement a robust **change management plan** and ensure **strategic communications** drive engagement and adoption amongst stakeholders. Throughout NAIRR's operation, it is critical provisioning and monitoring of devices are done on a continual and automated basis to prevent data leaks and unauthorized use.

Q3: How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Ethical and Trustworthiness evaluation criteria need to be part of the process when using the NAIRR and should be instantiated at the beginning of a research project. Anytime there is model development and training, there is a risk of biases being implicitly introduced. The NAIRR should consider how to mitigate risks from several types of biases, to include **data-driven bias** which can result from the data used to train the model, **confirmation bias** where models filter out relevant information from the user, as they think it is "noise", or **judgement bias**, where models can make determinations based on ambiguous data.

Any model developed needs to be evaluated to determine it meets the elements of Trustworthy AI. Rather than manual reviews, the NAIRR can develop an assurance framework to support the integrity of models and underlying data. Additionally, the NAIRR **could establish a quality review process which assesses certain properties of the data, models, and questions being answered**. For example, the review process could include acknowledging legal and

privacy restrictions, data consistency across population cohorts, data governance rules, and an assurance that datasets are representative of the issue the model is trying to solve.

Example: The NAIRR can draw from financial institutions that put algorithm risk into practice at scale. After the global financial crisis exposed the risks of inaccurate algorithm-driven models, the Federal Reserve and Office of the Comptroller of the Currency (OCC) issued the Supervisory Guidance on Model Risk Management (SR 11-7). It required identification and estimation of adverse consequences of inaccurate or misused models. SR 11-7 obligates users understand the limitations of models and avoid using models for uses other than originally intended. It mentions models should be validated regularly to ensure they are performing as expected.

Q4: What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Government

- **Data.gov:** GSA's Technology Transformation Services manages a repository of metadata, data.gov. The site stores 320+ datasets and resources such as tools, policies, and case-studies. Data.gov and other repositories such as Data USA could be leveraged to enhance data access for the NAIRR.
- **NIH Data Commons:** NIH ran a pilot for NIH Data Commons from 2017-2018. The pilot tested the best ways to build and implement a cloud-based platform designed to store, share, access and interact with data and other digital objects. For example, a tool called **FAIRshake** evaluates how easy / difficult it is to find, access, interoperate, and reuse digital objects. There are similar types of programs that exist, such as Dept. of Transportation's Safety Data Commons and USAID's Development Data Commons.
- **JAIC JCF:** JAIC is building an AI data and algorithm development platform called JCF. The NAIRR could establish a data exchange with JCF to bolster their data catalogue.

Academic

- **Colleges and Universities:** Schools across the country are standing up AI Institutes, and programs the NAIRR can tap into through partnerships. The University of California San Diego, along with five other universities, is trying to address the challenge of scaling across several areas including health, semiconductor chip design, robotics, and networks. In addition, numerous other university consortiums are looking to conduct research, such as the University of Chicago with the Digital Transformation Institute and Howard University, which is part of an Historically Black Colleges and Universities (HBCU) consortium developing curriculums and talent in quantum computing.

Private Sector

- **Computational Resources:** CSPs offer elastic compute and data storage capabilities ideal for a holistic computing ecosystem. By leveraging a hybrid cloud, the NAIRR can reap the benefits of cloud technology while still leveraging existing on-premises compute

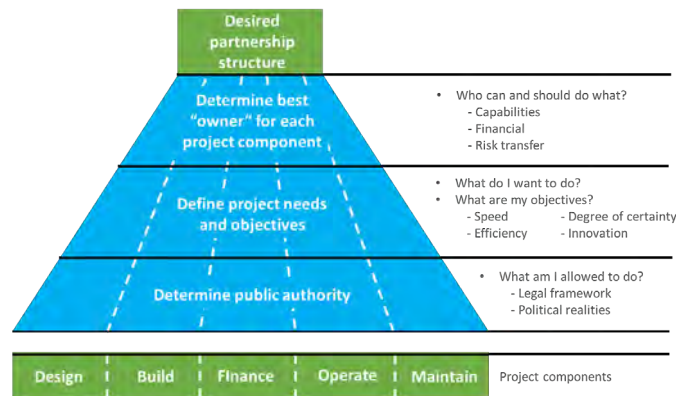
infrastructure owned by research hubs. In a hybrid model, NAIRR research hubs will have access to pre-built AI tools that support their research.

- **Private Sector Investments / Products:** The NAIRR can explore existing private sector assets to jump start the creation of the NAIRR with pre-configured capabilities. An example is [Deloitte’s Cortex AI](#), which draws from Deloitte use cases, solutions, and data, and acts as an accelerator to AI adoption by applying them to help clients devise intuitive solutions, scale adoption faster, and develop a competitive edge in their work.

Q5: What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public Private Partnerships (PPP) can be a powerful tool in tackling complex public service programs such as the NAIRR. PPPs can help government engage diverse stakeholders, control costs, and mitigate risks. However, **PPPs are not a one-size-fits-all**. In order to choose the right PPP model, government must determine the project components, public authority, goals, and the best “owner” for each phase from among the various service providers, vendors, system integrators, and startups that comprise the stakeholder ecosystem.

There are numerous models of PPP tailored to uses ranging from building highways to managing public goods such as airports or transportation infrastructure. PPPs have even been used to shape digital infrastructure in a similar way to what the NAIRR is trying to accomplish.



Source: Deloitte analysis

For example, NIH used PPP models to create NIH Data Commons, a cloud-based platform where researchers could store, share, access, and use digital files generated from biomedical research. The [NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability Initiative](#) widened the lens of partners, working with CSPs to allow 2,500 NIH-funded institutions to explore cloud and **ML capabilities to generate, analyze, and share data**. Recently, data has shown that these PPPs can work quickly and on sensitive data, creating a centralized, secure, and **cloud-enabled data platform** to analyze real-world COVID-19 patient data across government, hospital systems, and research institutions.

Q6: Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Data quality, management, and accessibility limitations

There are operational and technical barriers for the NAIRR to fulfill its goal of democratizing access to AI R&D. Operationally, success of the NAIRR will be heavily contingent on the platform's management of and access to quality datasets for users. As is often said in data science, *"garbage in garbage out"*, meaning bad data results in bad outputs. The NAIRR will need to gain / maintain the trust of users and develop mechanisms to encourage data sharing. For example, users from the scientific community may use journal standards to dictate whether submission of replication datasets / code are required. A solution is to develop data quality standards each research hub is encouraged to meet as a part of the NAIRR agreement. This could be incentivized through access to storage and / or additional cloud compute resources.

To tackle technical challenges, data cataloguing across the platform can help users quickly index large volumes of data to find relevant datasets. The NAIRR can offer templates and scripts to fix issues like formatting, missing values, and irregularities to help foster quality data. The NAIRR could survey data housing, assembly, cataloguing, and sharing practices across disciplines as a step toward establishing data quality standards that meet the needs of all user groups.

Cybersecurity Threats

There has been exponential growth in cyberattacks the past couple years, and for a program reliant on a decentralized governance model, it is imperative security mechanisms are implemented from bottom up to maintain program integrity. Platforms that house data and provide computational resources are charged with prioritizing confidentiality, integrity, and availability of data. For this, the NAIRR could leverage software to intelligently manage IdAM controls across the enterprise. The NAIRR could invest in tools to continuously monitor the network to increase visibility into network activity to detect potential security breaches, allowing the incident response team enough time to mitigate risks.

Organizational Culture

Organizations may be fearful or resistant to AI across their enterprise. Job displacement and ethical implications may dissuade users. The NAIRR can prepare R&D initiatives by emphasizing stakeholder inclusion, sensitivity to interests of data subjects and modeled outcomes. The NAIRR should incorporate a change management plan inclusive of all stakeholders and deliver strategic comms regularly to avoid misunderstandings or overlooked concerns and biases.

Diversity in AI

The blossom of AI has not yet been distributed in an equitable manner. Underrepresented groups have been locked out from both a talent (researchers and developers) and as consumers / beneficiaries of AI R&D. To help combat this, the NAIRR could consider mechanisms to prioritize allocated resources or research priorities that will benefit small startups and / or underrepresented institutions such as HBCUs.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Digital Diagnostics

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

October 1, 2021

2300 Oakdale Blvd.
Coralville, IA 52241

Eric Lander, PhD
Director
Office of Science & Technology Policy
Executive Office of the President
1650 Pennsylvania Avenue
Washington, D.C. 20504

Sethuraman Panchanathan, PhD,
Director
National Science Foundation
2415 Eisenhower Avenue
Alexandria, Virginia 22314

www.dxs.ai

Re: RFI Response: National AI Research Resource

Dear Dr. Lander and Dr. Panchanathan,

On behalf of Digital Diagnostics, a national healthcare innovator dedicated to health equity and access headquartered in Coralville, Iowa, we are pleased to offer our feedback on the “Request for Information: National AI Research Resource (NAIRR).” We believe that patient access to rigorously validate, U.S. Food and Drug Administration (FDA) authorized healthcare AI is an important tool to address social determinants of health and health equity. The autonomous AI service provided by our technology, IDx-DR (described by CPT code 92229), has transformed access to care for Medicare beneficiaries with diabetes and improved the ability of practitioners to prevent diabetes-related vision loss at the point-of-care. We offer below additional detail on our technology and our specific experience in the area of AI ethics and bias.

Digital Diagnostics’ Autonomous AI Technology

Digital Diagnostics, formerly IDx Technologies, is a pioneering autonomous AI diagnostics company on a mission to transform healthcare accessibility, quality, and affordability. Its flagship product IDx-DR (the service described by CPT 92229) is an autonomous AI system that was granted De Novo authorization by the U.S. Food & Drug Administration (FDA) to diagnose diabetic retinopathy and diabetic macular edema after rigorous FDA validation for safety and equity. The FDA determined that IDx-DR met the standards for “breakthrough device” designation in accordance with section 3051 of the 21st Century Cures Act. The technology has proven that intelligent diagnostic platforms can be used safely, efficiently, and equitably to improve patient outcomes.

Founded and led by Michael Abramoff, MD, PhD, a practicing physician and fellowship-trained retina specialist, Digital Diagnostics uses a patented biomarker-based approach to build autonomous AI, and a

rigorous, ethical framework for designing, developing, and deploying its AI.^{1,2,3} Driven to remove the health inequities in diabetes-caused vision loss and blindness by early detection and timely treatment, the IDx-DR system was designed to make the diabetic eye exam more accessible, especially for underserved populations, including racial and ethnic minorities and rural populations. The AI driven system performs the process of the diabetic eye exam at the point-of-care following similar cognitive processes as a highly trained eye care provider.

The IDx-DR technology diagnoses diabetic retinopathy and diabetic macular edema during a patient's routine diabetes management visit, closing a significant care gap. Thirty million people in the U.S. have diabetes, and diabetic retinopathy affects nearly 30 percent of diabetic patients.⁴ The disease is the most frequent cause of blindness among people ages 20-74 years old.⁵ In spite of the severity of diabetic retinopathy, nearly half of Medicare beneficiaries who have diabetes do not have an annual eye exam,⁶ and only 15 percent of people with diabetes have regular eye exams, as per the Standard of Diabetes Care from the American Diabetes Association (ADA).⁷ Diabetic retinopathy accounts for nearly \$500 million in total direct medical costs annually,⁸ much of which is due to late stage, undiagnosed disease.

Diabetic retinopathy and its concomitant visual loss and blindness is a major source of health disparities. Greater access to the diabetes eye exam is an important health equity issue, as significant health disparities exist in diabetes care and access to the diabetic eye exam. According to the U.S. Centers for Disease Control (CDC), the percentage of U.S. white adults with diagnosed diabetes (7.4 percent) is less than that of Black (12.1 percent) and Hispanic (12.7 percent) adults, and roughly half of American Indian /Alaska Native (15.1

¹ Abramoff, M. D., D. Tobey, and D. S. Char. "Lessons Learned About Autonomous Ai: Finding a Safe, Efficacious, and Ethical Path through the Development Process." *Am J Ophthalmol* 214, no. 1 (2020): 134-42. <https://dx.doi.org/10.1016/j.ajo.2020.02.022>.

² Char, Danton S., Michael D. Abramoff, and Chris Feudtner. "Identifying Ethical Considerations for Machine Learning Healthcare Applications." *The American Journal of Bioethics* 20, no. 11 (2020/11/01 2020): 7-17. <https://dx.doi.org/10.1080/15265161.2020.1819469>.

³ Abramoff, M.D., B. Cunningham, B. Patel, M.B. Eydelman, T. Leng, T. Sakamoto, R. M. Wolf, A.K. Manrai, J.M. Ko, and M.F. Chiang. "Foundational Considerations for Artificial Intelligence," *Ophthalmology* [in press] (2021). [https://www.aajournal.org/article/S0161-6420\(21\)00643-6/fulltext](https://www.aajournal.org/article/S0161-6420(21)00643-6/fulltext).

⁴ Lundeen EA, Wittenborn J, Benoit SR, Saaddine J. Disparities in Receipt of Eye Exams Among Medicare Part B Fee-for-Service Beneficiaries with Diabetes — United States, 2017. *MMWR Morb Mortal Wkly Rep* 2019;68:1020–1023. DOI: [http://dx.doi.org/10.15585/mmwr.mm6845a3external icon](http://dx.doi.org/10.15585/mmwr.mm6845a3external%20icon).

⁵ U.S. Centers for Disease Control & Prevention, "Vision Loss and Age" (accessed August 19, 2021), available at <https://www.cdc.gov/visionhealth/risk/age.htm>.

⁶ Lundeen EA, Wittenborn J, Benoit SR, Saaddine J. Disparities in Receipt of Eye Exams Among Medicare Part B Fee-for-Service Beneficiaries with Diabetes — United States, 2017. *MMWR Morb Mortal Wkly Rep* 2019;68:1020–1023. DOI: [http://dx.doi.org/10.15585/mmwr.mm6845a3external icon](http://dx.doi.org/10.15585/mmwr.mm6845a3external%20icon).

⁷ Benoit, S. et al. "Eye Care Utilization Among Insured People with Diabetes in the U.S., 2010-2014," *Diabetes Care* (March 2019) available at <https://pubmed.ncbi.nlm.nih.gov/30679304/>.

⁸ Rein et al. "The Economic Burden of Major Adult Visual Disorders in the United States," *JAMA Ophthalmology* (2006) available at <https://jamanetwork.com/journals/jamaophthalmology/fullarticle/418866>.

percent) adults.⁹ Although Medicare covers annual eye exams for all diabetic patients, approximately 55 percent of white Medicare patients with diabetes received an annual eye exam in 2017, while the prevalence of having an eye exam was lower for Black (48.9 percent) and Hispanic (48.2 percent) patients.¹⁰

The IDx-DR service provides a complete point-of-care service, including image acquisition, AI driven quality feedback, analysis and individualized, per patient results. As CMS noted in the preamble to the CY 2021 PFS final rule, “the AMA CPT Editorial Panel also created CPT 92229 (*imaging of retina for detection of monitoring of disease; with point-of-care automated analysis with diagnostic report; unilateral or bilateral*) for point-of-care automated analysis that uses innovative artificial intelligence technology to perform the interpretation of the eye exam, without requiring that an ophthalmologist interpret the results. CPT code 92229 can be used at an outpatient clinical setting and the artificial intelligence technology interprets the test instead of a remotely located ophthalmologist.”¹¹

Health Equity

The OSTP and NSF pose the question, “How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?” Our experience in AI ethics and bias may inform NAIRR’s efforts in this regard.

IDx-DR (described by CPT 92229) was designed specifically to create equitable access for diabetes patients in underserved and rural areas. The service has already been used to test thousands of patients across the U.S and has identified previously undiagnosed cases of disease that would have otherwise gone undetected and likely resulted in irreversible vision loss. For example, at the University Medical Center (UMC) in New Orleans, a predominantly Black patient population, there had been a backlog of over 800 patients with diabetes who had not received an eye exam. Wait times for a visit to an eye care provider in that area exceeded 4 months, leading many patients to give up seeking eye care. After adoption of the autonomous AI service in the diabetes clinic at the point of care, backlogs were entirely eliminated, and over 25% of the patients were identified with potentially blinding diabetic retinopathy that otherwise would have been entirely missed.^{12,13}

⁹ U.S. Centers for Disease Control, “Addressing Health Disparities in Diabetes,” *2017 Diabetes Score Card* (2017) available at <https://www.cdc.gov/diabetes/disparities.html>.

¹⁰ Lundeen EA, Wittenborn J, Benoit SR, Saaddine J. Disparities in Receipt of Eye Exams Among Medicare Part B Fee-for-Service Beneficiaries with Diabetes — United States, 2017. *MMWR Morb Mortal Wkly Rep* 2019;68:1020–1023. DOI: [http://dx.doi.org/10.15585/mmwr.mm6845a3external icon](http://dx.doi.org/10.15585/mmwr.mm6845a3external%20icon).

¹¹ CY 2021 Payment Policies under the Physician Fee Schedule and Other Changes to Part B Payment Policies, available at <https://public-inspection.federalregister.gov/2020-26815.pdf> (“2021 MPFS”), page 489.

¹² Harris, Richard, “FDA AI Challenge: How to Assess Safety and Effectiveness,” *National Public Radio* (April 19, 2019.) Available at <https://www.npr.org/sections/health-shots/2019/04/14/711775543/how-can-we-be-sure-artificial-intelligence-is-safe-for-medical-use>.

¹³ Company data. Digital Diagnostics. August 26, 2021.

In another clinical setting in Georgia, incorporating IDx-DR (CPT 92229) directly into the comprehensive diabetes management workflow increased compliance with the quality measure for evaluation of diabetic retinopathy from 16 to over 50 percent. Relatedly, the clinic almost tripled the number of previously undiagnosed vision related disease occurrences and was able to arrange immediate follow up with eye care specialists. And in Alabama, after implementation during the PHE at a Federally Qualified Health Center (FQHC) serving diverse and underserved beneficiaries with multiple chronic conditions, over 25% of the patients had potentially blinding diabetic retinopathy. Without the IDx-DR service, many of these patients would have gone on to develop costly vision threatening complications from diabetes. Finally, a cost effectiveness analysis from the patient perspective on IDx-DR installed at Johns Hopkins showed that autonomous AI for the diabetic eye exam reduces patients co-pay, compared to in-person eye care provider visit, thereby improving access for socioeconomically disadvantaged populations.¹⁴

In addition to the positive impact that increased access can have to advance health equity, we have also focused on how to design our clinical algorithms to eliminate bias and support equitable care, including the development of a rigorous ethical framework.^{15,16,17} Accordingly, when conceptualizing IDx-DR, we were concerned with the risk of racial, ethnic and other inappropriate bias during its entire lifecycle, namely, impact on patient benefit, algorithm design, algorithm training, AI validation, and populations in which it is implemented. Thus we chose the exact severity of diabetic retinopathy and macular edema diagnostic cutoffs, designed the AI to mitigate diagnostic and follow-up bias, and analyzed the distribution during the machine learning training phase.¹⁸ When designing the first ever preregistered pivotal trial for autonomous AI together with FDA, we created a novel endpoint for equity.¹⁹ This endpoint tested whether or not there was any race or ethnicity effect on the accuracy of the AI – the trial showed there was none. Finally, we chose to emphasize implementation in those populations where health disparities for diabetic retinopathy need to be addressed.²⁰

¹⁴ Wolf RM, Channa R, Abramoff MD, Lehmann HP. Cost-effectiveness of Autonomous Point-of-Care Diabetic Retinopathy Screening for Pediatric Patients With Diabetes. *JAMA Ophthalmol*. Oct 1 2020;138(10):1063-1069. doi:10.1001/jamaophthalmol.2020.3190

¹⁵ Abramoff, M.D., B. Cunningham, B. Patel, M.B. Eydeman, T. Leng, T. Sakamoto, R. M. Wolf, A.K. Manrai, J.M. Ko, and M.F. Chiang. "Foundational Considerations for Artificial Intelligence," *Ophthalmology* [in press] (2021). [https://www.aajournal.org/article/S0161-6420\(21\)00643-6/fulltext](https://www.aajournal.org/article/S0161-6420(21)00643-6/fulltext).

¹⁶ Char, Danton S., Michael D. Abramoff, and Chris Feudtner. "Identifying Ethical Considerations for Machine Learning Healthcare Applications." *The American Journal of Bioethics* 20, no. 11 (2020/11/01 2020): 7-17. <https://dx.doi.org/10.1080/15265161.2020.1819469>.

¹⁷ Abramoff, M. D., D. Tobey, and D. S. Char. "Lessons Learned About Autonomous Ai: Finding a Safe, Efficacious, and Ethical Path through the Development Process." *Am J Ophthalmol* 214, no. 1 (2020): 134-42. <https://dx.doi.org/10.1016/j.ajo.2020.02.022>.

¹⁸ Ibid.

¹⁹ Abramoff, M. D., D. Tobey, and D. S. Char. "Lessons Learned About Autonomous Ai: Finding a Safe, Efficacious, and Ethical Path through the Development Process." *Am J Ophthalmol* 214, no. 1 (2020): 134-42. <https://dx.doi.org/10.1016/j.ajo.2020.02.022>.

²⁰ Abramoff, M.D., B. Cunningham, B. Patel, M.B. Eydeman, T. Leng, T. Sakamoto, R. M. Wolf, A.K. Manrai, J.M. Ko, and M.F. Chiang. "Foundational Considerations for Artificial Intelligence," *Ophthalmology* [in press] (2021). [https://www.aajournal.org/article/S0161-6420\(21\)00643-6/fulltext](https://www.aajournal.org/article/S0161-6420(21)00643-6/fulltext).

Conclusion

Thank you for your efforts to provide feedback on the NAIRR. If you have any questions on our comments, please contact Juli Goldstein, MHS, Vice President, Government Affairs & Market Access at



Sincerely,

Michael Abramoff, MD, PhD
Founder and Executive Chairman
Digital Diagnostics, Inc.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Electronic Privacy Information Center (EPIC)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

NOMINATION OF THE ELECTRONIC PRIVACY INFORMATION CENTER

to the

White House Office of Science and Technology Policy and National Science Foundation

Regarding the

Implementation Plan for a National Artificial Intelligence Research Resource

October 1, 2021

EPIC submits the following comments to the White House Office of Science and Technology Policy and National Science Foundation on the Implementation Plan for a National Artificial Intelligence Research Resource (“NAIRR”).¹ EPIC urges the NAIRR Task Force to (1) devote significant resources to the robust assessment and preservation of privacy, civil rights, and civil liberties in the face of growing AI use; (2) provide regulators at the federal, state, and local levels with resources to ensure that civil rights and consumer protection laws are enforced against entities that deploy AI or automated decision-making systems; and (3) to limit partnerships with the private sector.

EPIC is a public interest research center in Washington, D.C. that was established in 1994 to focus public attention on emerging privacy and related human rights issues and to protect privacy, the First Amendment, and constitutional values.² EPIC has a long history of promoting transparency and accountability for information technology.³

EPIC has a particular interest in promoting algorithmic transparency and has consistently advocated for the adoption of the Universal Guidelines for AI (“UGAI”) to promote trustworthy and

¹White House Office of Sci. and Tech. Policy and Nat’l Sci. Found., *Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resources*, 86 Fed. Reg. 3908, <https://www.federalregister.gov/documents/2021/07/23/2021-15660/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>.

² EPIC, *About EPIC* (2019), <https://epic.org/epic/about.html>.

³ EPIC, *Algorithmic Transparency* (2018), <https://www.epic.org/algorithmic-transparency/>; EPIC, *Algorithms in the Criminal Justice System* (2018), <https://www.epic.org/algorithmic-transparency/crim-justice/>; Comments of EPIC, *Consumer Welfare Implications Associated with the Use of Algorithmic Decision Tools, Artificial Intelligence, and Predictive Analytics*, Fed. Trade Comm’n (Aug. 20, 2018), <https://epic.org/apa/comments/EPIC-FTC-Algorithmic-Transparency-Aug-20-2018.pdf>; Comments of EPIC, *Developing UNESCO’s Internet Universality Indicators: Help UNESCO Assess and Improve the Internet*, United Nations Educ., Sci. & Cultural Org. 5-6 (Mar. 15, 2018), [https://epic.org/internetuniversality/EPIC_UNESCO_Internet_Universality_Comment%20\(3\).pdf](https://epic.org/internetuniversality/EPIC_UNESCO_Internet_Universality_Comment%20(3).pdf).

careful adoption of algorithms.⁴ EPIC has advocated for transparency and accountability in the deployment of AI and algorithmic decision-making tools, litigating cases against the U.S. Department of Justice to compel production of documents regarding “evidence-based risk assessment tools”⁵ and against the U.S. Department of Homeland Security to produce documents about a program to assess the probability that an individual will commit a crime.⁶ In 2018, EPIC and leading scientific organizations petitioned the U.S. Office of Science and Technology Policy to solicit public input on U.S. Artificial Intelligence Policy.⁷ EPIC submitted comments urging the National Science Foundation to adopt the UGAI and to promote and enforce the UGAI in the funding, research, and deployment of U.S. AI systems.⁸ EPIC has also recently submitted comments to the National Security Commission on Artificial Intelligence, the U.S. Office of Science and Technology Policy, the European Commission, and the U.S. Office of Management and Budget urging robust regulation to protect individuals.⁹

In an effort to establish necessary consumer safeguards, EPIC has filed FTC complaints against HireVue,¹⁰ an employment screening company, and AirBnB,¹¹ the rental service that claims to assess risk in potential renters based on an opaque algorithm. EPIC has also filed a petition with

⁴See, e.g., EPIC v. DOJ, No. 18-5307 (D.C. Cir. settled Mar. 25, 2020), <https://epic.org/foia/doj/criminal-justice-algorithms/>; Comments of EPIC, *Intellectual Property Protection for Artificial Intelligence Innovation*, U.S. Patent and Trademark Office (Jan. 10, 2020), <https://epic.org/apa/comments/EPIC-USPTO-Jan2020.pdf>; Comments of EPIC, *HUD’s Implementation of the Fair Housing Act’s Disparate Impact Standard*, Dep’t of Hous. and Urban Dev. (Oct. 18, 2019), <https://epic.org/apa/comments/EPIC-HUD-Oct2019.pdf>; Testimony of EPIC, Mass. Joint Comm. on the Judiciary (Oct. 22, 2019), <https://epic.org/testimony/congress/EPIC-FacialRecognitionMoratorium-MA-Oct2019.pdf>; Statement of EPIC, *Industries of the Future*, U.S. Senate Comm. on Commerce, Sci. & Transp. (Jan. 15, 2020), <https://epic.org/testimony/congress/EPIC-SCOM-AI-Jan2020.pdf>; Comments of EPIC, *Request for Information: Big Data and the Future of Privacy*, Office of Sci. and Tech. Policy (Apr. 4, 2014), <https://epic.org/privacy/big-data/EPIC-OSTP-Big-Data.pdf>.

⁵ EPIC, *EPIC v. DOJ (Criminal Justice Algorithms)* (2020), <https://epic.org/foia/doj/criminal-justice-algorithms/>.

⁶ See *id.*; EPIC, *EPIC v. DHS (FAST Program)* (2018) <https://epic.org/foia/dhs/fast/>.

⁷ EPIC, *Petition to OSTP for Request for Information on Artificial Intelligence Policy* (July 4, 2018) <https://epic.org/privacy/ai/OSTP-AI-Petition.pdf>.

⁸ EPIC, *Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan*, Nat’l Sci. Found., 83 Fed. Reg. 48655 (Oct. 26, 2018), <https://epic.org/apa/comments/EPIC-Comments-NSF-AI-Strategic-Plan-2018.pdf>.

⁹ Comments of EPIC, *Solicitation of Written Comments by the National Security Commission on Artificial Intelligence*, 85 Fed. Reg. 32,055, Nat’l Sec. Comm’n on Artificial Intelligence (Sept. 30, 2020) <https://epic.org/apa/comments/EPIC-comments-to-NSCAI-093020.pdf>; Comments of EPIC, *Request for Comments on a Draft Memorandum to the Heads of Executive Departments and Agencies, “Guidance for Regulation of Artificial Intelligence Applications,”* 85 Fed. Reg. 1,825, Office of Management and Budget (Mar. 13, 2020) <https://epic.org/apa/comments/EPIC-OMB-AI-MAR2020.pdf>; Comments of EPIC, *Request for Feedback in Parallel with the White Paper on Fundamental Rights*, European Comm’n Fundamental Rights Policy Unit (May 29, 2020), <https://epic.org/apa/comments/EPIC-EU-Commission-AI-Comments-May2020.pdf>; Comments of EPIC, *Proposal for a legal act of the European Parliament and the Council laying down requirements for Artificial Intelligence*, European Comm’n (Sept. 10, 2020), <https://epic.org/apa/comments/EPIC-EU-Commission-AI-Sep2020.pdf>.

¹⁰ Complaint of EPIC, *In re HireVue* (Nov. 6, 2019), https://epic.org/privacy/ftc/hirevue/EPIC_FTC_HireVue_Complaint.pdf.

¹¹ Complaint of EPIC, *In re Airbnb* (Feb. 27, 2019), https://epic.org/privacy/ftc/airbnb/EPIC_FTC_Airbnb_Complaint_Feb2020.pdf.

the FTC for a rulemaking for AI in Commerce¹² and a complaint with the Attorney General for the District of Columbia concerning automated test proctoring tools.¹³ EPIC has also published the *AI Policy Sourcebook*, the first reference book on AI policy.¹⁴

In addition to EPIC’s responses to the specific questions posed about NAIRR, EPIC urges the Task Force to prioritize protecting civil liberties and civil rights over leading the world in AI deployment and development. EPIC warns the Task Force against incentivizing the deployment and development of AI simply for the sake of innovation and competition, given the threat to human rights posed by the increased collection of sensitive data and the use of inaccurate or discriminatory automated decision-making systems. It is these risks that led United Nations High Commissioner for Human Rights Michelle Bachelet to recently call on governments to “ban AI applications that cannot be operated in compliance with international human rights law and impose moratoriums on the sale and use of AI systems that carry a high risk for the enjoyment of human rights, unless and until adequate safeguards to protect human rights are in place.”¹⁵ Commissioner Bachelet also stressed the need for comprehensive data protection legislation in addition to a regulatory approach to AI that prioritizes protection of human rights. Commissioner Bachelet explained: “The risk of discrimination linked to AI-driven decisions—decisions that can change, define or damage human lives—is all too real. This is why there needs to be systematic assessment and monitoring of the effects of AI systems to identify and mitigate human rights risks.”¹⁶ Accordingly, the Task Force should take this opportunity to drive *thoughtful* and *responsible* development and deployment of AI.

EPIC Answers to Specific OSTP/NIST Questions):

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

The Task Force should prioritize the assessment of “privacy, civil rights, and civil liberties requirements associated with the NAIRR” above other “capabilities and services.” The protection of privacy, civil rights, and civil liberties should guide the other capabilities and services listed, including making government data sets available as part of the NAIRR.¹⁷

¹² EPIC, *Petition for Rulemaking Concerning Use of Artificial Intelligence in Commerce* (Feb. 3, 2020) <https://epic.org/privacy/ftc/ai/EPIC-FTC-AI-Petition.pdf>.

¹³ EPIC, *In re Online Test Proctoring Companies* (2021), <https://epic.org/privacy/dccppa/online-test-proctoring/>.

¹⁴ *EPIC AI Policy Sourcebook 2020* (2020), <https://epic.org/bookstore/ai2020/>.

¹⁵ United Nations Human Rights Council, *The Right to Privacy in the Digital Age*, A/HRC/48/31 (Sept. 13, 2021), <https://epic.org/UN-AI-Moratorium-Call-And-Report.pdf>.

¹⁶ *Id.*

¹⁷ White House Office of Sci. and Tech. Policy and Nat’l Sci. Found., *supra* note 1, at 39,082 (“As outlined in § 5106(b) of Public Law 116-283, the implementation roadmap developed by the Task Force should include the following:

- A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;
- B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:
 - i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and
 - ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

NAIRR should provide resources for companies and government entities developing AI to safeguard privacy and to address issues of racial and gender bias, fairness, civil rights, transparency, and accountability. For resources NAIRR will create for developers, this means recommending prohibitions on certain uses of AI and the collection of certain types of data as well as guidance on data minimization and strict restrictions on data sharing and selling. The Task Force should also create an accountability toolkit that developers should be encouraged to implement throughout their design and deployment process—including independent mandatory validation studies, civil rights impact assessments, and algorithmic audits. In addition to this, the NAIRR should provide guidance for government agencies procuring AI or automated decision-making systems to be purposeful and transparent about what they procure.

In some countries, including Canada, certain AI systems used in public contexts must undergo Algorithmic Impact Assessments that evaluate the risks posed by an individual system based on the sensitivity of data used, design attributes, and relation to areas designated as requiring additional considerations and protections.¹⁸ For example, the Canadian tool prompts an entity deploying an AI system to evaluate the stakes of the decisions that the system in question makes, the vulnerability of subjects, and whether the system constitutes a predictive risk assessment.¹⁹ The tool also allows for multiple answer options and a detailed explanation of responses. The Canadian assessment also requires detailing the downstream processes of an AI system, including (i) whether the system will only be used to assist a decision-maker; (ii) whether the system will be replacing a decision that would otherwise be made by a human; (iii) whether the system will be replacing human judgment; (iv) whether the system is being used by the same entity that developed it; (v) consideration and explanation of both economic and environmental impacts; and (vi) consideration of the sensitivity of data collected.²⁰ NAIRR should incorporate a resource that addresses these questions for both developers and contractors and a template for transparency concerning

-
- C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;
 - D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;
 - E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;
 - F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;
 - G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;
 - H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; and
 - I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.”)

¹⁸ Canada Digital Servs., *Algorithmic Impact Assessment* (2021), <https://open.canada.ca/aia-eia-js/?lang=en>.

¹⁹ *Id.*

²⁰ *Id.*

developers, factors used, data sources, and more. EPIC particularly urges the Task Force to consult the resources referenced below in Question 4 that provide guidance on government procurement of new technologies.²¹

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

EPIC recommends that the Task Force consider the following resources and principles as building blocks for the implementation of the NAIRR:

- Public Voice Coalition, [*Universal Guidelines for AI*](#)
 - Right to Transparency.
 - Right to Human Determination.
 - Identification Obligation.
 - Fairness Obligation.
 - Assessment and Accountability Obligation.
 - Accuracy, Reliability, and Validity Obligations.
 - Data Quality Obligation.
 - Public Safety Obligation.
 - Cybersecurity Obligation.
 - Prohibition on Secret Profiling.
 - Prohibition on Unitary Scoring.
 - Termination Obligation.²²
- Organisation of Economic Cooperation and Development, [*OECD AI Principles*](#)
 - Inclusive growth, sustainable development and well-being.
 - Human-centered values and fairness.
 - Transparency and explainability.
 - Robustness, security and safety.
 - Accountability.²³
- Rashida Richardson, [*Best Practices for Government Procurement of Data-Driven Technologies*](#) (2021)
- World Economic Forum, [*AI Procurement in a Box: AI Government Procurement Guidelines*](#) (2020)
- Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, Madeleine Clare Elish, Jacob Metcalf, [*Assembling Accountability: Algorithmic Impact Assessment for the Public Interest*](#), Data & Society (2021)
- Mona Sloane, [*The Algorithmic Auditing Trap*](#), OneZero (Mar. 17, 2021)
- Rebecca Kelly Slaughter, Janice Kopec, & Mohamad Batal, [*Algorithms and Economic Justice: A Taxonomy of Harms and a Path Forward for the Federal Trade Commission \(Info. Soc. Project & Yale J. Law & Tech.\)*](#) (2021).

²¹ Rashida Richardson, *Best Practices for Government Procurement of Data-Driven Technologies* (2021); World Econ. Forum, *AI Procurement in a Box: AI Government Procurement Guidelines* (2020).

²² *Id.*

²³ *Recommendation of the Council on Artificial Intelligence*, OECD (May 21, 2019), legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449.

- United Nations High Comm’r for Human Rights, [*The Right to Privacy in the Digital Age*](#) (2021)
- Ben Green & Amba Kak, [*The False Comfort of Human Oversight as an Antidote to A.I. Harm*](#), Slate (June 15, 2021)
- Amba Kak & Rashida Richardson, [*Suspect Development Systems: Databasing Marginality and Enforcing Discipline*](#), 55 U. Mich. J.L. Reform (forthcoming 2022)
- Rashida Richardson, [*Racial Segregation and the Data-Driven Society: How Our Failure to Reckon with Root Causes Perpetuates Separate and Unequal Realities*](#), [*Berkeley Technology Law Journal*](#), Vol. 36, No. 3 (forthcoming 2022).

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public-private partnerships should play an extremely limited role in the NAIRR in order to maximize independence and prioritization of civil liberties and rights. If public-private partnerships are necessary, NAIRR should set rigorous restrictions on the influence of companies involved and on the access those companies have to government data sets.

Conclusion

EPIC looks forward to engaging further with the NAIRR and urges the Task Force to help make the NAIRR a trusted resource for the responsible development and deployment of AI. The NAIRR should prioritize privacy, civil rights, and civil liberties and create useful resources that can operationalize these principles and bring AI deployment into compliance with consumer protection and civil rights laws.

Respectfully Submitted,



Ben Winters
EPIC Counsel

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Engine

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

September 1, 2021

National AI Research Resource Task Force
Attn: Ms. Wendy Wigen, NCO, NITRD Program
2415 Eisenhower Avenue
Alexandria, Virginia 22314

RFI Response: National AI Research Resource (NAIRR)

Dear Members of the National AI Research Resource Task Force,

Engine is a non-profit technology policy, research, and advocacy organization that bridges the gap between policymakers and startups. Engine works with government and a community of thousands of high-technology, growth-oriented startups across the nation to support the development of technology entrepreneurship through economic research, policy analysis, and advocacy on local and national issues. We appreciate the opportunity to comment on the development of NAIRR. AI research and development can be prohibitively expensive and out of reach for many startups operating on bootstrap budgets.¹ The establishment of NAIRR has the opportunity to foster competition and innovation by creating opportunities for startups to work in the AI space without incurring all of the R&D costs associated with AI development.

Question 1, Item A on goals for establishment and sustainment of NAIRR and metrics for success

The creation of NAIRR is a tremendous opportunity to encourage innovation and boost competition in AI. AI R&D is an incredibly data-heavy and resource-intensive undertaking, which makes it inaccessible for many startups. According to an Engine report from earlier this year, the average seed-stage startup is working with about \$55,000 in monthly capital.² Outside of places like Silicon Valley, this figure drops below \$50,000 per month.³ Because of the expense associated with AI R&D and their tight budgets, some startups in the AI space rely on grants and other small chunks of funding to carry them through to their next formal funding round.⁴

To ensure that NAIRR can boost innovation and competition, the Task Force must count startups as beneficiaries and make NAIRR available to them. This goal of startup access is aligned with the vision of the program as articulated by the authors of the law. When the House passed the National

¹ See, e.g., Ivy Nguyen, “Could data costs kill your AI startup?,” VentureBeat, Nov. 10, 2018, <https://venturebeat.com/2018/11/10/could-data-costs-kill-your-ai-startup/>, (conceptually walking through the data value chain).

² “the State of the Startup Ecosystem,” Engine, Apr. 2021, <https://static1.squarespace.com/static/571681753c44d835a440c8b5/t/60819983b7f8be1a2a99972d/1619106194054/Te+State+of+the+Startup+Ecosystem.pdf>.

³ Ibid.

⁴ Edward Graham, “#StartupsEverywhere: Owings Mills, MD.,” Engine, Oct. 2, 2020, <https://www.engine.is/news/startupseverywhere-owings-mills-md-sofia-labs>.

AI Research Resource Task Force Act as part of the National Defense Authorization Act for Fiscal Year 2021⁵ and again when the House passed the conferenced version of the NDAA⁶, Rep. Anna Eshoo (D-Calif.) included private companies when describing beneficiaries of NAIRR: “The national AI research cloud expands access so that American universities and companies can participate in AI R&D.”

Question 1, Item C on a model for governance and oversight

If startups are going to be a beneficiary of NAIRR, it’s critical that members of the startup ecosystem should have a seat at the table as the Task Force considers questions about strategic direction and programmatic decisions. This will also make the program more dynamic and responsive to practical issues as it continues, ensuring it’s continued utility for startups and the promotion of innovation. The Task Force should consider seeking out the perspectives of ecosystem support organizations (ESOs)—such as incubators, accelerators, and startup-related non-profits—to get an overarching view of the startup perspective.

Question 1, Item E on barriers to the dissemination and use of high-quality government data sets

Many government datasets at present are inaccessible to many stakeholders that could derive value from them. The Task Force should seek to avoid this issue as it establishes a roadmap for NAIRR. Practical access should be facilitated for stakeholders seeking to use NAIRR, including by startups for commercial uses. This should include proactive outreach to startups and members of the startup ecosystem (ESOs) and the provision of workshops regarding the use of NAIRR.

Question 5 on the role of public-private partnerships in NAIRR

As laid out above, startups should have access to NAIRR for commercial use. There is a long history of private entities, including businesses, making use of government data for things like market research, investment planning, and other business functions. Indeed, the OPEN Government Data Act specifically enumerates the facilitation of data use for non-government entities including businesses.⁷ NAIRR should be no different.

Public-Private partnerships can also play an important role in the reach, effectiveness, and success of NAIRR. The government should conduct proactive outreach to stakeholders, including startups, that could benefit from NAIRR. This outreach should leverage members of the startup ecosystem like ESOs that startups turn to for information and resources. Many government agencies, from the Small Business Administration, to the Patent and Trademark Office, to the Census Bureau, conduct outreach, hold webinars, and host workshops aimed at increasing the use of their programs.

⁵ 166 Cong. Rec. H3501 (Jul. 20, 2020) (statement of Rep. Eshoo)

⁶ 166 Cong. Rec. H6932 (Dec. 8, 2020) (statement of Rep. Eshoo)

⁷ “Public Law 115–435: OPEN Government Data Act.” (132 Stat. 5536; Jan. 14, 2019)

Question 6 on democratizing access to AI R&D

As explained in response to Question 1, Item E and in response to Question 5, above, startups can only be benefitted by government resources and programming if they know about it. In order for NAIRR to be effective in fostering innovation and promoting competition in the AI space, the government must facilitate access for startups, including through proactive outreach. This could include promoting NAIRR at startup-facing events like conferences and hackathons as well as introducing ecosystem support organizations across the country to NAIRR so they can educate the startups in their respective networks. Consulting with ecosystem support organizations as the Task Force prepares a roadmap would likely generate additional engagement opportunities.

Thank you for your work on developing the roadmap for NAIRR and the opportunity to submit input.

Sincerely,

Engine

Engine
700 Pennsylvania Ave. SE
Washington, DC 20003

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

The Enterprise Neurosystem

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Introduction

The Enterprise Neurosystem is an open source community of Fortune 500 companies, enterprise technology vendors, and academia. It is ultimately designed to address the fundamental challenges of large-scale AI infrastructure to protect our global ecosphere. The founding academic institutions include Stanford SLAC, Harvard Analytics, and UC Berkeley Data-I. Participating firms include America Movil, Equinix, Fiducia | AI, IBM Research, Intel, Kove, PerceptiLabs, Verizon Media/Yahoo!, Red Hat, and others.

The Challenge

Climate change, energy security, agricultural challenges, political unrest, and mass migration pose immediate threats to our existence and the planet's health as a whole. The exponential growth of technology provides us the means to address these challenges, but only if we, as a nation, can build an enabling infrastructure to manage them effectively. As a result, we will unlock innovations as heterogeneous systems merge and reveal cross-domain patterns in data, algorithms, technology, and ultimately new discoveries.

The Enterprise Neurosystem community sees the National Artificial Intelligence Research Initiative as a means to build the requisite wide-ranging AI infrastructure that will identify and circumvent environmental threats, provide advanced insight regarding human migration, integrate climate-adaptive agriculture from local to global scales, and integratively advance fusion energy, nuclear security, and nuclear medicine while guaranteeing privacy and security. The mission here is to promote a sustainably accelerated human advancement in synchrony with the extended longevity of life on the planet.

Such a planetary scale objective begins with smaller developmental steps at the intersection of the scientific community and the Fortune 500 enterprise market. Development at this junction will foster a mid-tier environment for the baseline AI technology that will progress to the end state objective. It extends well beyond academic research, yielding a continuous stream of commercially viable architectures that will act as an autonomous and self-aware AI ecosystem that further advances the objective. This cannot be overstated, the rapid spin-off productization enhances the next technological development, and this compounded multiplicative effect is exactly the root of the Kurzweilian exponential technological growth law.

A handful of major companies cannot foster and harvest the rich intellectual capital in the US research and technology sectors. This Intellectual capital is best served by a quintessentially US culture that encourages and rewards ambitious entrepreneurialism and creative expression. In

this way, our national innovation culture provides ideal raw creativity to incubate the creative solutions needed for the US and the world's challenges. We feel that a well-designed and forward-minded National AI Research Initiative would provide the tools and the connectivity to exponentially accelerate our nation's use of AI to solve the globe's largest problems.

Groupthink runs contrary to American innovation culture because it reduces the necessary exploration of design space. Therefore, open communities of creative problem-solvers more frequently sow new and innovative ideas. This is clear from the acquisition cycle of startups by the more resilient larger corporations. The startups are bleeding-edge and collectively cast a wider net of innovation. And the successful ideas are then picked up by major companies which are more adept at scaling the ideas to broad application.

In the field of AI, however, innovation always begins from a very data intensive position. New ideas require enormous high-performance computing (HPC) and data resources unavailable to garage-level startups. We propose the **"National AI Initiative"** to create a national infrastructure that enables and encourages this development and provides the tools and platforms that foster discoveries and faster time to market.

The Solution

The community has noted the preponderance of AI models being deployed in the enterprise. They currently exist as bespoke solutions that lack deep integration with other models or domains, thus precluding the innovation that only arises from the cross-fertilization of collective findings. It seems clear that a connective data fabric and a centralized cross-correlation model are required for a national AI infrastructure that aligns with the network's long-term objective to monitor the planet and take real-time actions as needed. This connectivity will involve the emerging swarm of Edge IoT devices, data sources, and AI models, with a core interpretive model and recommendation engine.

A distributed infrastructure paradigm shift needs to take place. For example, instead of exporting data into one giant repository, standardized and composable feature extraction methods can be applied to improve efficiencies in distributed model training and enforcement of confidentiality. Data should only be moved on demand or by way of interpretive metadata or interface layer, one that both increases security and reduces expense and extraneous network traffic. Metadata can be delivered in a tiered fashion to reduce latency and shorten the time to action. Data generalities will only give way to finer granularity based on user requirements, permissions, and authentication. For instance, data sampling techniques that include anonymization through hashing can lead to smaller data sources maintained in their respective silos in a federated model. Data and related features are shared only on an as-needed and as-granted basis.

Such a national data fabric infrastructure with a curated and multi-tiered security system to authenticate users and enable targeted data sharing would be a requisite primary focus of this new architecture. Independent AI-powered security instances will travel the network to scan and authenticate new users and data sources. In a non-intrusive manner, round-the-clock penetration testing will be enabled, and remote decryption key monitoring and related pattern analysis will be implemented. Instead of granting access to the entire network, focusing on Layer 7 application connectivity based on mTLS and Zero Trust Network Architecture (ZTNA) will isolate and reduce the impact of intrusions to a single application.

Although Public Cloud Offerings have demonstrated industries' acceptance and hunger for cloud-based HPC services, the newly emerging Edge computational and AI paradigms create novel security challenges as these resources are decentralized. The more advanced hardware is not currently available in Public Cloud Infrastructure. At the same time, users can already see tremendous benefit from a Federated model that exposes innovations in a secure sandbox. In such a sandbox, both scientific researchers and industry innovators could explore tailored hardware, even work together to co-design new technology for their specific research or market needs. A federated model would also allow data to remain resident in individual silos, fostering a safe but rich collaborative outcome via its tiered metadata capabilities.

Production Architectures

Institutions Contribute AI Models to a Shared Registry



Illustration 1 - Shared Registry with decentralized data

The basic premise is that the various institutions and individuals will use a shared registry of common AI models. This registry becomes a catalog of models, with different versions trained for specific purposes. These models can be transferred, then used and reused in each edge location, running on-premise data to tune these collectively pre-trained models. Both bespoke and general core models will exist in the shared registry, available to be shared across locations without moving the resident data. Retention of training based on the individual locations will be regulated as per the access/authentication controls and the intent of the original data producer, thus maintaining the desired security paradigm and privacy controls.

Collaboratively create AI Models across institutions

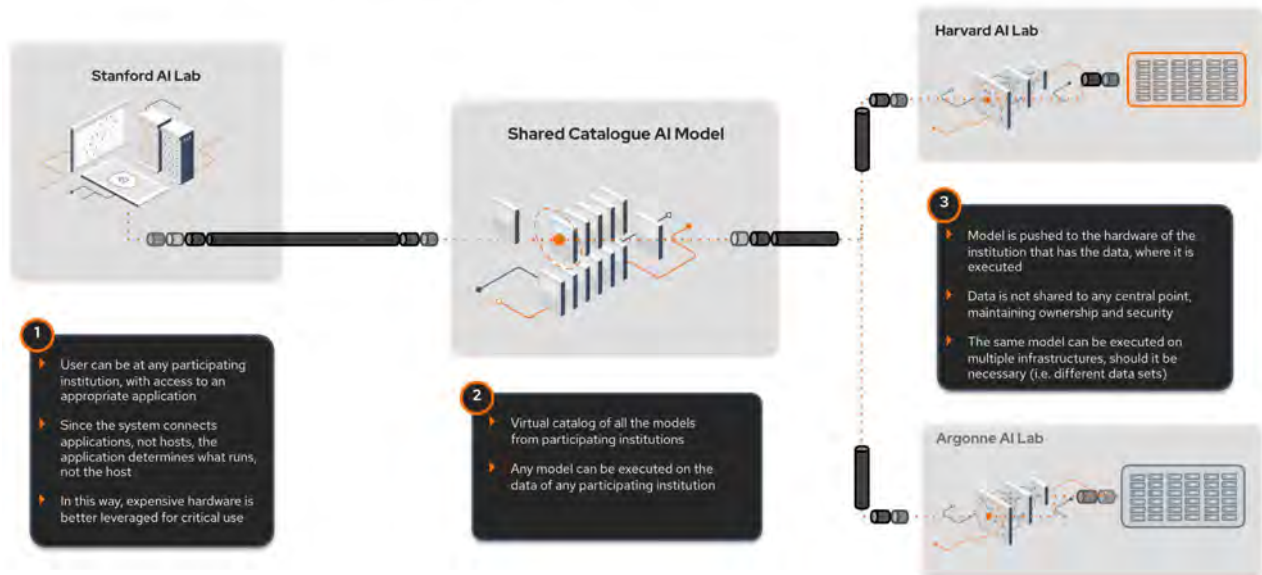


Illustration 2 - Secure Collaboration Diagram

The shared infrastructure enables users at one institution to access models at another institution and execute those models on data of any participating institution. In this way, users are not limited to only running models at their institution or data gathered only at their institution. Further, the utility of models is greatly enhanced because they can be used against many different data sets. A true shared model is created.

Managing a Shared Catalog for Enhanced Collaboration

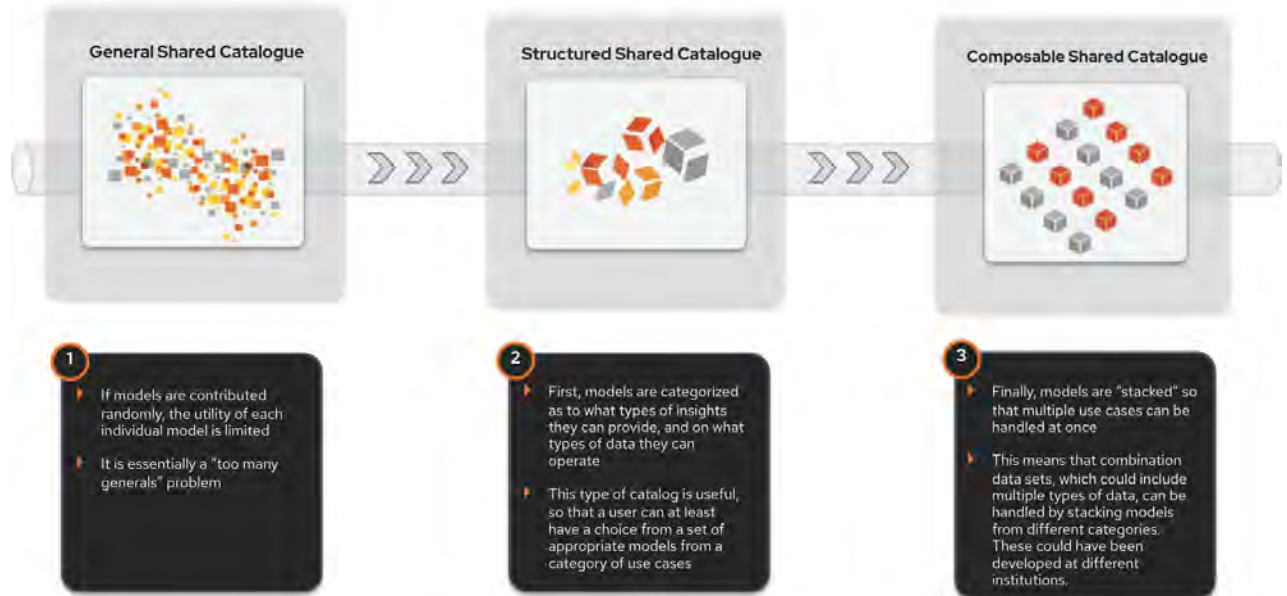
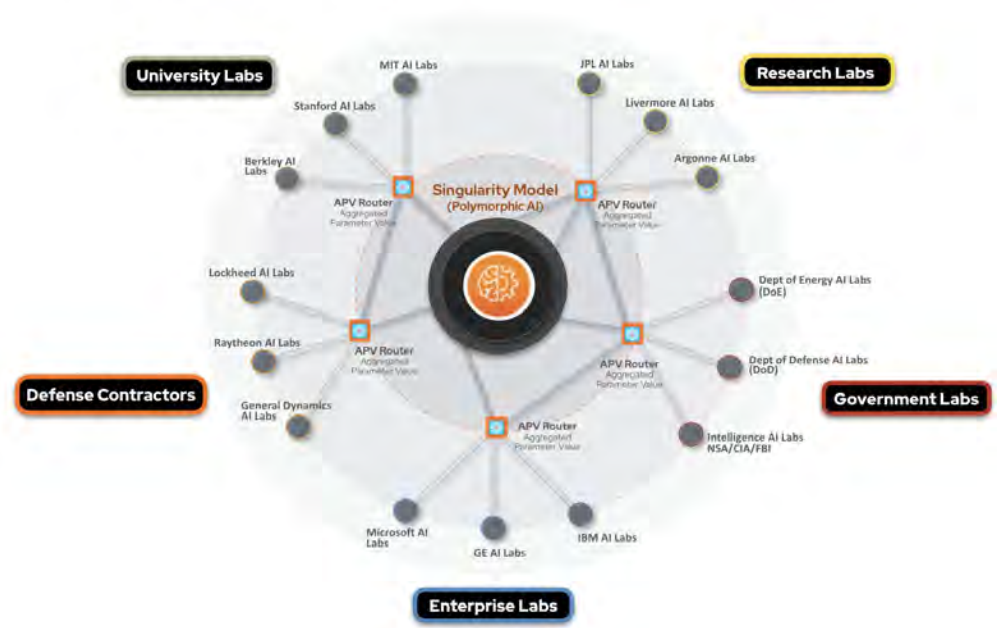


Illustration 3 - Managing the shared repository/catalog

Since models are dependent on the types of data they act upon, it's crucial to organize models by the kinds of data they can process. Once completed, the models can be "stacked" so a stack of 2 appropriate models can execute that data composed of two other types of data. This is the essence of the "composable shared catalog."

Proposed Singularity Architecture

- Architecture Fundamentals**
- Models are trained locally at each institution
 - Trained model parameters are aggregated and sent to update and enhance the "singularity model" that is the basis of the experiment being run
 - Actual data is never shared with any central point, or any other institution
 - The next iteration will have updated parameters that include the benefit of everyone's training runs
 - An overarching singularity model can be used for dedicated cross-correlation and deeper pattern analysis across all models
 - This leads to deeper and more accurate insights, real-time corrective actions, and targeted recommendations for users



For this framework, singularity refers to the greatest synergy between human and artificial intelligence. It is humanity's greatest ally in the realm of scientific discovery, industrial advances, and the fight against environmental threats. A central cross-correlation engine that conducts ongoing pattern analysis, security reviews, and bias health checks will cross all the activities within the various resources. It is an opt-in model, and not every classified lab function will require (or be given) access to this resource. But in the case of the national labs, a real-time cross-correlation function can advance science with deeper insights and course correction during experimental campaigns. The complete body of lab research then becomes a national AI resource. The enterprise and AI startups can help propel their intellectual property forward with new feature discovery, digital twin testing, and rapid time to market.

One primary intelligence can analyze all incoming and historical data, or it can be a series of singularities that are purpose-built for different fields of endeavor. In any event, having this capability gives a far more profound understanding of our environment and all its possibilities.

Compartmentalizing Security to Application Specific Framework

Moving from VPN to Zero-Trust Network Architecture Security




Illustration 5 - Application Specific Framework with Security Isolation

Since we are using a shared infrastructure model, data security must be handled more granularly than in traditional approaches. It is not sufficient to have a “single access point” that grants access to all models and all data – such a system would be vulnerable to a single security breach giving access to everything (as is shown in the traditional architecture on the left).

On the right side of the diagram, the more modern approach shows how individual application access controls can be applied so that having access to one application does not give access to all resources. These can even be timed and/or revocable access tokens. A single security breach would not result in more than one application being compromised.

Questions Section:

1. What options should the Task Force consider for roadmap elements A through I, and why?

In terms of RFI elements A through C, it seems clear that an open community approach to the ownership and maintenance of this resource would be optimal. Essentially, a federated compute, storage, and AI development environment composed of multiple labs and data centers is connected on a highly secure and lossless framework with a small group of dedicated paid resources to manage the common infrastructure instance. This funded management team would coordinate architectural upgrades and responses to technical issues.

The related hardware and software resources would be donated and shared by the various organizations and participants in a federated network architecture. A governance model for this resource could be based on open community principles via the Linux Foundation or similar community frameworks. Furthermore, given element D, it requires a multi-tiered approach to data access and provenance tracking through Distributed ledger-based technology. All domains will likely benefit significantly from HPC-enabled metadata and digital twin creation. This motivates an infrastructure that enables a unified framework among a diversity of components, from the IoT Edge to the HPC Core. Such heterogeneity will be critical to the success of this endeavor, and several platforms and emerging techniques can be combined to help address this requirement.

Commonly known required components:

- Open and closed data sets that help participants enable model training capability to be identified, built, and curated. Metadata frameworks would be created to increase efficiency and help navigate issues of privacy and bias.
- A generalized digital twin environment that supports discovery and data set generation.
- A user-permission and authentication mechanism.
- A platform to create, distribute and manage AI models, with capabilities including ground-up development, pre-built models, pipeline workflow, lifecycle management, and drift correction.
- A hardware environment including all necessary resources, including storage, processors (GPU/TPU/IPU, x86, ARM), networking components, and software infrastructure platforms (Kubernetes, container management, databases, memory optimization, etc.).

Unique elements would include:

- A flexible and dynamic resource federation model that is based on an automated data fabric that extends across all elements. This helps private sector firms gain access to

and directly build better hardware and software via a shared space that is co-designed with public sector emerging needs.

- A marketplace that enables both open and private science and industry, with models available to any organization. In essence, a library of templated architectures and base pre-trained AI models for related research purposes that ultimately enable cross-correlation of models, incorporating relevant heterogeneous data sources and sensors that can generate pattern analysis in real-time to enable deeper insights and rapid course correction.
- A library of composable transformation and featurization layers that aid user adoption of the templated architectures and enable security and provenance tracking from the beginning of the development cycle.
- A software-defined memory allocation and virtualization framework that eliminates bottlenecks for large-scale AI workloads and optimally capitalizes on both on-premise and distributed data stores.
- A tiered security architecture that includes:
 - Ongoing 24/7 penetration testing with non-intrusive security scans.
 - Distributed ledger-based user/resource authentication.
 - Layer 7 integration strategy (intrusions relegated to a single application).
 - Biometric authentication and resource isolation within the federated network for sensitive workstreams.
 - Monitor decryption keys for data provenance, audit, and automated discovery of abuse patterns.
 - Independent software instances that autonomously assign themselves to assess and scan newly federated hardware.
- Open source fairness and bias management tools, provided by a dedicated community ethics group. This team will apply operational parameters to data sources, metadata layers, and cross-correlation engines to maintain systemic neutrality. This includes a tight focus on privacy and the protection of individual rights.

2. What capabilities and services provided through the NAIRR should be prioritized?

- Curated and openly available data sets would be the top priority in our estimation.
- Various users and institutions can access openly available models in a registry.
- Readily available hardware infrastructure for model training, openly available to all member institutions. Beyond training requirements, the use of more expensive chipsets (GPUs, TPUs, IPUs, FPGAs, etc.) should be evaluated from a price/performance perspective against standard chip architectures (x86, etc.) in terms of production lab use. This will lead to a balanced workload approach, cost savings, and co-designed hardware and algorithm compositions.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

FABRIC Testbed

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

FABRIC Testbed Response to the Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource (NAIRR)

Submitted by:

FABRIC Testbed:

Ilya Baldin, RENCi

James Griffioen, University of Kentucky

Tom Lehman, Virnao

Inder Monga, DoE ESnet

Anita Nikolich, University of Illinois-Urbana Champaign

Paul Ruth, RENCi

KC Wang, Clemson University

1. What options should the Task Force consider for any of roadmap elements A through I above, and why?

Roadmap element A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success:

- **Goal 1: Infrastructure for AI compute**
 - Emphases: extensible, distributed, federation, clear/single entry/portal, security, open to all (while allowing diversified features, privileges, and governance)
 - Metrics: Capacity, sustainability, number and diversity of stakeholders, clarity, and low barrier of use
- **Goal 2: Infrastructure for AI Data**
 - Emphases: extensible distributed, flexible movement, federation, clear/single entry/portal, security, open to all (while allowing diversified features, privileges, and governance)
 - Metrics: Diversity, number and size of accessible data sets, clarity, and low barrier of use
- **Goal 3: Infrastructure and tools for sharing and templates for AI Workflow**
 - Emphases: best practice, programmable templates supporting flexible use of any NAIRR resources, community & learning portal.
 - Metrics: Number, types, and quality of templates, use counts of templates
- **Goal 4: Risk Framework and governance processes**
 - Emphasis: balance of open access and advanced, strategic mission uses

- Metrics: coverage of uses, efficiency of approval, supervision, and auditing of proper use of infrastructure

2. Which capabilities and services provided through the NAIRR should be prioritized?

A NAIRR can be the enabler to bring together (or federate) disparate, individual system testbeds that specialize in singular technical aspects such as cybersecurity, cloud computing, autonomous vehicles, 5G, network, etc., into a larger, overarching AI testbed ecosystem that simulates how AI will work across organizational and technology boundaries. Although discipline-specific testbeds host experiments containing elements of machine learning, combining them under a larger umbrella - a testbed of testbeds - enables a holistic, system-wide view of AI. Forming an overarching “AI Testbed of Testbeds” based on existing new infrastructure such as FABRIC will enable more meaningful analyses of the ways in which results from AI decision making serve as input into and influence upon the decisions made by other systems. This is particularly important as machine learning data from one system feeds other systems; variations in this input are important to test in the structured manner a testbed provides.

Testing AI systems for resilience and safety is currently done in a number of verticals that remain isolated from each other. AI for Medical, Vehicles, IoT devices, Adversarial AI and AI-driven Networks are all tested in varying, specialized testbeds. However, the dependent interconnections of AI systems across these verticals can benefit from a larger testbed ecosystem that enhances repeatability and provides auditability and oversight. An overarching AI Testbed of Testbeds would encourage reproducibility and enable data sharing across AI domains since access to the Testbed ecosystem is expected to be more democratized than current efforts that are restricted to academia or government. An AI testbed will also enable realistic adversarial machine learning and AI environments, the results of which enhance AI resilience and trust.

A National AI Testbed of Testbeds could interconnect Internationally to similar nation-scale systems and AI testbeds around the world. For example, results of ML research based on experimental sensors in one smart city or edge testbed should be easily repeatable using several other smart cities’ testbeds. Any differences in outcomes may indicate opportunities to tune AI/ML algorithms and/or optimize smart cities.

A NAIRR Testbed enables experiments that explore open research questions such as:

Adversarial AI and AI Security. 3000+ papers on attacking (and to a lesser extent defending) ML models have been published in the past ~5 years. While very few of these attacks are detected in the real world, a place for experimentation at scale and across multiple technologies is important in order to build resilient AI systems that are protected from attacks. Additionally, no singular testbed exists to do vulnerability assessments of AI models. A testbed environment in which models exist on multiple types of devices will assist security researchers in assessing new types of security defenses which must be constructed. There is very little knowledge among practitioners on how to secure AI systems.

Developing Machine Learning and AI development best practices. Requirements about how to fairly develop and secure AI models are lacking, as are risk management standards and even widely agreed upon best practices. A testbed is an ideal, sandboxed environment in which to explore emerging work in this area.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Ethical and responsible AI R&D must follow a process which is open and auditable and should be conducted on an independently operated infrastructure open to anyone with the proper credentials. NAIRR can help to establish this national testing resource. AI systems research and development requires testing - ideally on a controlled and well-instrumented testbed under simulated real-world conditions. Current nation-scale systems testbeds focus on individual technologies such as 5G, cloud computing or networking. AI-enabled systems will require testing across multiple areas of computer science. A NAIRR encourages research to be performed on a neutral, nation-wide infrastructure which requires documentation of the training data set provenance, information flow provenance as well as experiment results. This will serve to enhance the transparency of algorithm construction and performance.

Meaningful AI research depends on data. A lot of data. Storing and sharing machine learning training data remains a challenge, especially in medical domains, although they are increasingly using Federated Learning and other privacy preserving methods. However, there is an opportunity during the development process by using a testbed to ensure data contains elements of fairness. Requiring documentation such as data set "nutrition labels" in the course of testing will prompt researchers to consider whether the data set is truly representative of all populations. This will also serve to force transparency in algorithm design and enhance auditability as data set provenance is documented. As part of the approval process for projects using a testbed, testbed owners can encourage researchers to address questionable data set

issues if the training data sets and data are made transparent. Part of an onboarding process can include a forcing function to have researchers answer questions about ethical sourcing, vetting, repeatability, etc.

An AI Testbed of Testbeds can uniquely serve as a conduit for an AI Data Exchange or perform a Data Sets as a Service function for those who don't have easy access to large amounts of labelled or unlabeled data. Since resilient and trustworthy AI systems should be tested, the testbed facility is the perfect matchmaker for those who have data and those who need it.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

The NSF community has accumulated a substantial amount of practical experience in building and operating large distributed testbeds for their research communities (GENI, DETER, NSF Clouds, PAWRs, SAGE, FABRIC). Lessons learned include architecting, deploying large distributed infrastructures, proper federated identity management, security procedures, control, measurement software, dealing with operational and experimenter data and, importantly, methods of operating these infrastructure by inter-institutional teams - lessons that can be transferred onto NAIRR to make it a sustainable, multi-institutional effort.

These testbeds contain important resources that can be tied together to support a variety of applied AI research, which when coupled to NAIRR infrastructure would create unique opportunities for cross-disciplinary experimentation in applying AI techniques to multiple technologies and research domains with significant impact on our society - 5G, autonomous infrastructure management, autonomous vehicles, cyber-security. Such resources, which a standalone NAIRR infrastructure would otherwise have to replicate at great cost.

Specifically, the NSF-funded FABRIC Testbed represents an important element of this ecosystem intended to tie together the many pieces, constructing a continuum of computing from large centralized resources to small edge resources and dedicated instruments, with an intelligent, programmable and stateful network in-between. As many data intensive disciplines seeking novel AI applications transition from the idea of processing data in situ to operating on streaming data as it is produced by instruments and sensors, FABRIC provides a unique platform on which such ideas can be tested at scale.

FABRIC with its broad reach to a wide variety of research communities also has significant convening powers through workshops and other outreach activities, providing a platform to discuss different types of research relevant to NAIRR goals.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Technology companies and content providers conduct much of the same machine learning research in parallel with academia but with a larger amount of both hardware resources and data. An increasingly large percentage of accepted papers at the top machine learning and AI conferences are from industry. But there has been little incentive to work together. Industry researchers have ample resources, while academic researchers struggle for funds. However, both industry and academia lack a nationwide multi-disciplinary testbed with resources dedicated to conducting machine learning and AI research. The benefit to industry is increased access to highly skilled researchers, while the benefit to academia is access to hardware resources and data. Both sides benefit from access to large scale, multi-disciplinary resources on which to conduct research. Several successful NSF programs such as PAWR and the recent Fairness in AI have been formed by cooperation between agencies and industry.

A joint, multi-agency effort, including NSF, DoE, DoD, DHS, CISA, FDA, FAA, and multiple others, would ensure that efforts at creating an AI Testbed aren't duplicated across Federal agencies.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Insider-only Data Access: Many of the advances in AI/ML/DL will be based on access to rich, extensive, data sets that, today, are often held and managed by companies or government agencies where access is restricted to approved members of the organization. In order to gain access to these data sets, researchers must often become a member of the organization (e.g., summer employment) or work with someone within the company or agency. These types of selective engagements do not scale and do not facilitate democratized access to the data. While industry is known to have some of the most extensive data sets, there will also be data sets needed for research in niche areas. These will often be collected by smaller sets of researchers who, historically, have been hesitant to share their research data, or will only share their data after they have studied it and some amount of time has passed. Such practices could also significantly limit data democratization efforts.

Screening Research Intentions: Good work is being done by organizations with research goals and intentions that might not align with those of the data holders. For example, non-profits, human rights organizations, and citizen scientists are all exploring questions related to fairness, particularly as it relates to AI/ML models, and may experience challenges accessing the data they need to carry out their research.

Resource Limitations: AI/ML research often requires massive data sets. The resources needed to store, transport, and process these data sets can be so burdensome that some portions of the research community are not able to participate. While some of these issues can be addressed with increased funding, others, such as local network infrastructure may be difficult to address for any number of reasons.

Legacy Access Methods: Existing data access techniques (e.g., remote file services and repositories, web services, network transport protocols, etc.) are not designed to provide efficient and precise access to the massive data sets needed by AI/ML/DL. Many AI/ML/DL systems are designed to pull in data sets from centralized or cloud storage and perform local computations, oftentimes pulling in the entire data set when only a small portion of the data is needed. This leads to slow data retrieval times and many redundant (local) copies of the data. Intelligent access mechanisms and protocols are needed to strategically pull in precisely the data that is needed. Testbeds such as FABRIC could be used to help develop these types of intelligent access mechanisms.

Viewing Data as Infrastructure: Data is problematic in terms of access to training data, data labeling, data sharing, etc. One cannot assume that the data can be moved to the processing. Current approaches to working with data, particularly data that is intentionally distributed and must remain where it is, need to be rethought. For example, Federated Learning and other new data and model pipelines as an architecture need a place to experiment at scale and testbeds such as FABRIC offer such a place. Storage of the data, the AI models, the results should be funded by a Federal entity.

Converged AI Research: As AI research rapidly evolves in the scope and scale of data, several critical considerations have emerged. Data are often ingested from multiple sources, and ensuring effective access, transfer, staging, and quality control of such data requires high quality tools, infrastructure, and best practices. A nationally shared testbed like FABRIC is not just about the state-of-the-art infrastructure it provides. Even more important is its provision of a shared infrastructure where **templates of curated toolsets, data, and computing infrastructure** can be created based on best practices and shared with researchers with diverse backgrounds and interests, so that all can begin developing their cutting edge AI research on a **common, high-quality, validated foundation**. The template approach is also an opportunity to bootstrap and broaden **transparency and ethical considerations** in AI research by publishing, sharing, and teaching such tools and methods amongst AI researchers to inspect their own AI data and solutions. Sharing of data and software tools has been a longstanding practice in many research communities. NSF XSEDE, for example, has facilitated sharing of High Performance Computing (HPC) tools in a national community. What FABRIC is offering is one step beyond the sharing of tools, instead, providing a complete environment with tools, data, instructions, and programming interfaces for developing new AI solutions.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Center for Security and Emerging Technology

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

September 27, 2021

RFI Response: National AI Research Resource
White House Office of Science and Technology Policy and National Science Foundation
86 FR 39081; Document Number 2021-15660

The Center for Security and Emerging Technology (CSET) offers the following submission for consideration by the Office of Science and Technology Policy and the National Science Foundation. OSTP and NSF requested information to support the work of the National Artificial Intelligence Research Resource (NAIRR) Task Force, which has been directed by Congress to develop an implementation roadmap for a shared research infrastructure that would provide Artificial Intelligence (AI) researchers and students across scientific disciplines with access to computational resources, high-quality data, educational tools, and user support.

Our submission recommends:

- **The NAIRR should facilitate access to compute through NSF's new CloudBank program.**
- **The NAIRR should provide computational resources for undergraduate and graduate students in the U.S. through cloud access beyond what is freely available through Google's Colab or similar industry resources.**
- **The NAIRR should provide computational capacity comparable to the compute used to train the private sector's largest models for specific research by PhD researchers.**
- **The Task Force should engage with the Department of Commerce and state commerce departments to determine an appropriate mechanism that allows entrepreneurs to access NAIRR compute resources.**
- **The NAIRR should create a data consortium, potentially through a public-private partnership (PPP), that cultivates, enriches, and maintains datasets for public use.**
- **The NAIRR should facilitate the sharing of data brought by researchers to CloudBank and enable sharing by default.**

202100261.A) Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success:

The NAIRR should nurture the development of AI to advance U.S. economic competitiveness and national security in areas that are underserved by the current ecosystem. In particular, the NAIRR should support research irrespective of potential commercial implications; support research into the safety, security, privacy, and equity of AI systems; and provide resources to underserved communities of developers.

Since the NAIRR is not empowered to change market incentives, policymakers should determine what it can *provide* or *aggregate* for the AI research community. Data, computational capacity, and talent--the people who create algorithms and design ML systems--are three main inputs for AI development; capital underpins access to all three. NAIRR's greatest impact will be to

improve the quality of, and access to, these inputs for a wide collection of actors that the market does not adequately support.

In order to establish the NAIRR’s structure, oversight, and metrics for success, CSET conducted a gap analysis of the AI development ecosystem using the matrix below. We considered the matrix for four different prototypical AI researchers who might benefit from the NAIRR: an undergraduate student, university PhD faculty or graduate student researcher, a PhD faculty with exceptional computational requirements, and an entrepreneur.¹ Actors in the horizontal axis shape the ecosystem of AI research using the inputs on the vertical axis. For example, the federal government provides computational capacity (infrastructure) to some PhD and graduate student researchers through the 18 National AI Research Institutes funded by NSF. **The analysis determined that undergraduate students and entrepreneurs rely on free computational resources provided by industry and that all four typical users suffer from a paucity of usable data.**

Figure 1. Matrix of U.S. R&D actors and inputs²

		Actors					
		Federal Government	State & Local Government	Industry	Philanthropy	Academia	International Partners
Inputs	Funding						
	Human Capital						
	Infrastructure						
	Policy & Regulation						

1.C) A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources.

The Networking and Information Technology Research and Development Program (NITRD) oversees Cloudbank, a program that negotiates rates and finance of cloud compute for researchers. The Task Force should leverage the NITRD's existing Cloudbank stewardship by giving it the responsibility and necessary authorizations to manage NAIRR resources. Placing the NAIRR under NITRD, and subsequently CloudBank, will speed the rollout of NAIRR cloud resources, facilitate access by NSF-funded researchers already eligible for CloudBank, and eliminate unnecessary duplication of similar efforts.

¹ Completed framework analysis available upon request.

² Melissa Flagg and Paul Harris, "System Re-engineering: A New Policy Framework for the American R&D System in a Changed World" (Center for Security and Emerging Technology, September 2020). <https://doi.org/10.51593/20200050>

1.D) **Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.**

Compute:

The NAIRR should facilitate cloud access through NSF’s new CloudBank program. Rather than creating another duplicative management structure, the Task Force could bolster the CloudBank program, providing accounts and credits to students and researchers across the U.S.

The NAIRR should provide computational resources for undergraduate and graduate students in the U.S. through cloud access beyond what is freely available through Google’s Colab or similar industry resources.

The NAIRR can achieve this aim by creating a student interface on CloudBank. NSF CloudBank could afford all student accounts a basic budget for compute and build a mechanism that lets them purchase more time, lets their university credit their accounts, and lets them receive scholarships for time on the platform. If adequately liberalized, such a mechanism would allow for outside funding to support students’ learning and research on the platform. This system would also enable students to augment classes both at their institution and through MOOCs by performing computationally intensive exercises that may have been previously out of reach. This structure would also subsidize university costs to teach AI classes, since students could access such a platform for their work rather than the university paying for compute. Students should also be able to file patent applications based on research conducted on NAIRR cloud resources, not subject to Bayh-Dole regulations that currently guide patent protections for university faculty. Such a comprehensive system would nurture self-guided learning, innovation, and democratize computational access to students attending universities without the means to afford similar resources.

Student activity should be co-signed by faculty to avoid misuse of computational resources, like mining cryptocurrencies. Misuse of the cloud resources could result in the loss of federal aid to the student, just as those charged with possession of illegal substances lose access to FAFSA support.

At the other end of the spectrum, **some researchers may need computational capacity comparable to the compute used to train the private sector’s largest models--the NAIRR should meet that need.** Owing to the expense of offering such computational capacity, applications for these extra-large models should be evaluated by the NSF as it does with other research grants. Providing enough compute may be easy, if a little costly. Meeting or in some cases exceeding the largest private sector models would cost as much as several million dollars which pales in comparison to other high profile research efforts like the National Ignition Facility (\$3.5B) or CERN (\$23B). High uptake of the NAIRR’s cloud compute will lead to increasing costs, but this should be considered an indicator of success—a signal that NAIRR resources are supporting previously unmet demand. Private sector cloud solutions including AWS, Microsoft Azure for researchers, and Google TPU Research Cloud already offer

researchers access to high performance computing (HPC). As the financial intermediary, CloudBank can negotiate more competitive rates for these services as usage increases overtime.

Separately, CloudBank could use existing private-public infrastructure to quickly increase access to high performance compute. For example, Frontera, an NSF-supported supercomputer located at University of Texas at Austin, already provides access to HPC to some AI researchers and could be incorporated into CloudBank. Incorporating existing public infrastructure into CloudBank, like Frontera, would create a centralized marketplace for researchers and students to access compute. We support a proposal by Dr. Vince Kellen, CIO of UC San Diego and a co-PI for CloudBank, that would increase cloud compute resources. According to Kellen, many researchers and academic institutions maintain on-premises compute capacity; how much is unknown. Some of this capacity is likely underutilized, as there can be downtime between experiments and computational cycles. Dr. Kellen identifies this as an opportunity to create a marketplace for compute: a single, unified mechanism where academic institutions can list and lease their excess capacity. This would maximize the use of existing computational capacity, by allowing researchers at other institutions to conduct research remotely. If such a system proved successful, the high performance computing facilities of the national labs may similarly benefit from such a market—though appropriate guardrails would be necessary given the value of such assets.

The Task Force should engage with the Department of Commerce and state commerce departments to determine an appropriate mechanism that allows entrepreneurs to access NAIRR compute resources. Entrepreneurs are a particularly difficult group to reach with resources. Unmoored by commitments to academic or other public institutions, entrepreneurs also represent the riskiest group to serve. Ethical questions answered by IRBs in academia are unanswered in independent business ventures, so a careful administrative process to evaluate individuals and their actions of NAIRR cloud resources is necessary. The Department of Commerce and state commerce departments may find it easier to negotiate access if a system of ethics review boards are put in place.³

But, it is also necessary to provide this access. Start-ups rely on capital, either from lenders or private equity, to purchase compute and data access. Although free data and compute are available, adequate amounts of either can be expensive. This pushes entrepreneurs into a position of needing a proof of concept before seeking investments or loans. Such a scheme could be paid for overtime if a small equity share of companies developed on the platform is held by an independent trust, which uses the proceeds of equity sales or revenue earned from dividends to pay for more computational resources later on or reimburse past expenses. Many start-ups may fail; therefore, any such revenue would only subsidize some of the operational costs.

Data:

The NAIRR should create a data consortium, potentially through a public-private partnership (PPP), that cultivates, enriches, and maintains datasets for public use.

³ James E. Baker, "Ethics and Artificial Intelligence" (Center for Security and Emerging Technology, April 2021). <https://doi.org/10.51593/20190022>

Procuring new data is time consuming and capital intensive. Well-financed companies can dedicate large amounts of resources to shaping data into usable formats for AI and machine learning research. Students, start-ups, and academic researchers cannot. The NAIRR must prioritize curating high-quality datasets for users—without it, computational resources may achieve relatively limited new discoveries as many researchers scour well-trodden data sources for new insights. The data consortium should own the following responsibilities:

1. **Identifying** datasets that are relevant to researchers’ needs and represent a diverse set of topics. If AI research is to enable the development and growth of other sectors, data that reflects those fields is crucial to researchers’ success.
2. **Acquiring** the data by scraping it from public-domain internet sources, buying it from vendors, manually recording data points contained in narrative sources such as news media, etc.
3. **Moving** the data from different endpoints (API, FTP, web downloads, etc.) into a computing environment where it can be processed and stored.
4. **Structuring and enriching** the data so that it has a consistent, meaningful format and can be efficiently sorted and queried. This may involve converting the data to a standard format, translating it from other languages, adding metadata, or developing and applying taxonomies and classification models. Outside specialists, such as annotators, technical experts, and translators, may be involved.
5. **Integrating** the data with other relevant datasets by mapping out common features, creating unique cross-connecting identifiers and fixing duplications and ambiguities. This “data fusion” process is a prerequisite for analyses involving more than one type of information (that is, most useful analyses).
6. **Validating** the data to ensure it reflects “ground truth” and has been processed correctly. This requires careful judgment and clear, consistent procedures.
7. **Documenting** the data so that others can understand what it is and how to use it.
8. **Hosting** the data through an online repository, likely data.gov or directly on the PPP’s website. The data should be readily available on the CloudBank (or other selected) computing platform.
9. **Maintaining** the data as facts, user needs, vendor terms, etc. change over time.

The NAIRR should integrate the data.gov website with the existing CloudBank platform and the proposed student access (above). Integrating the two offerings may increase the use of government data.

The NAIRR should facilitate the sharing of data brought by researchers to CloudBank and enable sharing by default. Researchers should be able to limit the availability of their data with a provided justification, up to and including keeping it all private, but sharing data has many potential benefits. To facilitate this data sharing, CloudBank would need to host storage space for data and allow user accounts with appropriate permissions to access the data. Besides introducing replicability to research findings and publications, sharing data can facilitate other research, encourage the development of new benchmarks on shared data sets, and break down systemic moats that protect some particular researchers who benefit from close ties to particular

government agencies or data sources. Although all data sharing is ultimately the subject of agreement between the original party and the data provider, access to NAIRR cloud resources could facilitate more collaborative behavior.

1.E) An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

While there are sometimes legitimate reasons for protecting government data, the government should identify non-sensitive data and share more of it. The NAIRR can help government agencies publish their data for use by researchers where feasible.

The Task Force should request an Executive Order or administrative rule requiring government agencies to 1) identify data that can be published without overriding classification or PII concerns, 2) fund NITRD to hire new FTEs able to conduct data quality analysis and standardization/cleaning of data sets in preparation for publication, and 3) require agencies work with these new FTEs to facilitate the publication of government data. NITRD will require new funding to hire these FTEs. As an alternative to new FTEs under NITRD, the Task Force could coordinate with the Presidential Innovation Fellows program, the U.S. Digital Service, and the U.S. Digital Corps. The end result will be new staff to help agencies clean and publish data for use by researchers of all stripes, not just AI researchers on CloudBank. The resulting curated data should be uploaded to data.gov. The mandate for agencies to work with new NITRD staff and publish those cleaned data will make available data which is currently held in an unusable format.

1.F) An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls.

The NAIRR should follow existing best practices for access management control utilized by private sector cloud service providers. Student's university accounts should be authorized by an email address from an accredited institution of higher education and require a cosigner from the institution.

Maintaining the integrity of datasets used on the NAIRR platform is an imperative. At a minimum, datasets uploaded to the NAIRR platform should be automatically hashed using a reputable hashing algorithm not known to be compromised (e.g., SHA256). Hashing datasets of data.gov by default, which is outside the purview of the Task Force, would also help authenticate the veracity of data before it is utilized. Though tedious, hashing datasets will prevent malign actors from tampering with data used to train models and conduct research. This step will help defend against data poisoning attacks and model backdoors.⁴

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

⁴ Andrew Lohn, "Poison in the Well: Securing the Shared Resources of Machine Learning" (Center for Security and Emerging Technology, June 2021). <https://doi.org/10.51593/2020CA013>

The first priority should be NAIRR access for entrepreneurs and start-ups. These individuals and firms can increase market competition, drive innovation, and challenge incumbents. Disruptive start-ups are a politically neutral way of challenging the market dominance of a few firms and can lead to entirely new applications of AI. Moreover, these individuals are caught between not having enough compute and data to produce a viable concept, which stops them from attracting private equity or receiving bank loans.

The second priority for the NAIRR should be to create full-time staff positions that identify, curate, and facilitate the publication of government datasets. The positions can be organized under NITRD, Presidential Innovation Fellows, U.S. Digital Service, or U.S. Digital Corps; no matter which agency presides over the effort, data publication must be their top priority. There is no mandate and few incentives for government agencies to share their data nor are they adequately staffed to prioritize the curation of these data sets for public use. Rather than each agency spinning up their own staff to do this work, the NAIRR should create specialized teams of data scientists with the authority to work with across agencies to execute data sharing efforts. All data that qualifies for cleaning, formatting, and publication should be hosted on data.gov.

Notably absent from these priorities is access for academic researchers. Their absence reflects our conclusion that reaching different user groups requires different levels of intentionality. If compute and data are made available to entrepreneurs and students on CloudBank, academic researchers with access to research funds will be able to easily purchase compute on the same platform. Supporting access for entrepreneurs and facilitating data publication by the government will take far more work and resources than helping academic researchers get access, and thus they must receive higher prioritization when evaluating resource distribution. If CloudBank is sufficiently liberalized, academic researchers will utilize the platform as they see fit.

4.) What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Compute:

The National Science Foundation supports CloudBank, an intermediary for NSF-funded researchers seeking to perform their research in the cloud. Though currently limited to 150 specific principal investigators, the platform could easily expand to other NSF-funded researchers and students. Because CloudBank negotiates financing for computational resources, more users will increase CloudBank's negotiating power.

- CloudBank's authorizations and responsibilities should expand to a wider set of researchers including students. This expansion should include infrastructure for account verification and management.
- CloudBank will need additional funding to allow student accounts to access these resources. Current CloudBank funding facilitates the purchase of cloud compute from companies with money from NSF-funded researchers' grants.

Another NSF-funded cloud program, Exploring Clouds for Acceleration of Science (E-CAS), also offers cloud services to 6 research programs. E-CAS is currently a smaller initiative than

CloudBank, but could serve as another pathway to providing cloud compute access to researchers.⁵

Frontera, currently the world’s fifth fastest high performance computer, provides compute to AI-focused researchers.⁶ Many national labs host high performance computers, including many of the 18 NSF-funded AI Research Institutes.⁷ If the Task Force chose to follow a “market-maker” model of supporting AI research—in addition to supporting CloudBank, then the NAIRR could host a single unified registration system to reserve time on high performance computers. This approach would encounter significant organizational barriers, but would leverage existing resources to support the most computationally intensive research.

However, access to compute must be simplified for researchers, particularly students. They should make a single application for compute resources and, after approval, where they actually receive compute should be transparent to them. By creating a one-stop shop for compute under the auspices of the NAIRR, researchers can avoid having to apply for compute resources across many programs.

Data Aggregation:

The federal government’s data.gov website publishes data from federal, state, and local government agencies. Data.gov should be a resource for researchers accessing the NAIRR. A website plug-in for data.gov on the NAIRR platform would suffice.

The private sector and research community compile public datasets with zeal.⁸ NAIRR would do well to link these resources on the NAIRR platform.

⁵ <https://internet2.edu/cloud/exploring-clouds-for-acceleration-of-science/>

⁶ <https://nsf.gov/cise/ai.jsp>

⁷ <https://nsf.gov/cise/ai.jsp>

⁸ [Google Cloud Public Datasets](#) | [AWS Public Datasets](#) | [Wikipedia Dumps](#) | [Kaggle](#) | [Harvard Dataverse](#) | [FigShare](#) |

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Google

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



Via electronic submission at <http://www.regulations.gov>

Subject: Google Comment Regarding the Office of Science and Technology Policy and National Science Foundation NAIRR Task Force Request for Information

Reference: 86 FR 39081, Document Number 2021-15660

October 1, 2021

Google welcomes the opportunity to provide comments in response to the **“Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource (NAIRR)”** issued jointly by the National Science Foundation and the Office of Science and Technology Policy.

Artificial intelligence (AI) will continue to have a significant, positive impact on society—for example through enabling improvements in healthcare, education, business, and government, and helping to address key challenges like climate change. The implementation and uptake of AI also raises questions about governance, privacy, security, access, economic opportunity, and dual use. We strongly support establishing a NAIRR, and share the government’s goal to make AI access more equitable through this resource. The NAIRR is a great opportunity to support increased access for a diverse range of researchers to critical AI and cloud resources such as storage, compute, databases, networking, data analytics, AI services and collaboration tools.

In developing a NAIRR implementation plan, we encourage the Task Force to focus on leveraging the existing and unique capabilities of US academic institutions, government agencies, and industry to further enhance US competitiveness.

1. What options should the Task Force consider for any of the roadmap elements, and why?

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success

Suggested goal: Provide efficient and expanded access to resources (including storage, compute, databases, networking, data analytics, AI services and collaboration tools) for US AI researchers and practitioners, from all sectors

(including academia; federal, state, and local governments; and private sector), and facilitate cross-sectoral and interdisciplinary projects. Researcher access should be prioritized based on an analysis of the public benefit of the proposed work and cost to provide such benefits. Special consideration should be made for how government and academic researchers may leverage the resource. Metrics for this goal:

- For funded proposals, time between a researcher's request for resources and access to resources
- # of overall NAIRR users, and # of active users
- Demographics of NAIRR users, including race/gender, users from specific sectors, regions, and populations, including government researchers, academics (with disaggregated data on academics from minority serving and emerging research institutions), and startups. We especially encourage use of the NAIRR to facilitate US government AI research, including interagency collaboration on AI and research at national labs.
- # of publications resulting from NAIRR use
- Project impact, measured through some combination of: # and size of follow on projects/funding, qualitative impact analyses, citations, impact factor of journal articles, awards
- # of open-source projects / contributions related to NAIRR use, and use of these contributions by the broader community
- # of interdisciplinary (i.e. led by chemists, biologists, humanities scholars) and cross-sectoral projects and publications, and publication impact
- Applications resulting from NAIRR use, with qualitative evaluation of their impact
- # of unique public/private data sets made available and used by researchers, and quality of de-identification of data in these data sets
- # of machine learning models/architectures made available and used by researchers
- Overall carbon impact and renewable energy savings, as well as carbon impact per project (goal should be carbon neutral to start, and move toward 100% renewable energy use for all compute, with yearly targets for renewable use)

Suggested goal: Workforce development -- train more AI researchers and practitioners, specifically in use of cloud resources. One of the main barriers to successful use of the cloud - as identified by federal programs such as CloudBank and STRIDES -- is that researchers lack training on how to access the cloud, how credits work, and how to appropriately budget for its use. Such training should be built into NAIRR implementation. In addition to training and support, NAIRR administrators should establish a formal process/platform to collect feedback from participating researchers and share with private sector training organizers, where relevant. Metrics for this goal:

- # and features of researchers and research administrators trained in specific aspects of cloud use, and correlation between this training and further use of the NAIRR (i.e. metrics described under the first goal above)
- Features can include demographics, regional representation, non-AI specialist researchers,
- Feedback received from users, and response to this feedback through program improvements.

Suggested goal: Facilitate AI applications discovery and implementation in the short- and long-term. The NAIRR should track use of resources for AI applications with high social impact -- for example, forecasting climate disaster likelihood with better accuracy and resolution, and precision medicine (i.e., biological datasets to design medical treatment tailored to individuals). It should also monitor use for short- (1-2 years from application), medium- (3-5 years from application) and long-term (6-10 years from application) projects, and evaluate whether there is an optimal ratio of projects falling into these buckets.

Metrics for this goal:

- # of milestones reached or applications discovered related to [NAE Grand Challenges](#)
- # of publications from NAIRR use that contribute to projects across diverse fields, and # of citations of these publications
- Time horizon for projects and breakdown across NAIRR portfolio between short-, medium-, and long-term projects.

B. Ownership & Administration

We can envision multiple models of government ownership and administration of the NAIRR, including establishment of a new FFRDC to oversee, a public-private partnership, or ownership by a particular agency with oversight and governance through an interagency body and external input. We encourage the Task Force to explore what governance models have been most effective for past public-private and interagency efforts. Wherever this resource is housed, we urge the NSF, in close coordination with OSTP, to ensure alignment with the multitude of other related NSF activities, such as the HPC Consortium, Open OCN, National Compute Research, AI Institutes, and Cloudbank. Similarly, other agencies with Cloud pilot programs should ensure that the NAIRR does not duplicate these efforts (e.g. the NIH STRIDES program) and instead are able to potentially form a more networked approach.

C. Governance & Oversight

Regardless of which agency or organization manages the NAIRR, we recommend that the National Science and Technology Council oversee its operations and ensure that all relevant agencies are encouraged to contribute data and tools to the resource, that government agencies are focused on using the resource for priority research and application areas, and that NAIRR implementation takes into account lessons learned from other federal cloud programs. We also recommend that the National AI Advisory Council create a subcommittee to provide oversight and feedback to the program, and that members of this subcommittee are chosen to represent the diverse user base for the NAIRR.

The NAIRR should deploy a two-tiered review process to select users for access to NAIRR compute. Most users would apply for “base-level” access to the resource, which would be granted through a light review process. Proposals requiring compute over a certain threshold level would be subjected to a more rigorous merit review process, which would prioritize use cases with the greatest potential to contribute to public knowledge and the public good. *Note that we have not set a threshold for “base level” versus higher levels of access, and would encourage the NAIRR task force to develop a specific recommendation here after consultation with possible applicants.* Other criteria for this merit review should include (but are not limited to): scientific merit; qualifications of the team to undertake the proposed research; whether the proposal is from an underrepresented user group (this could be measured by type of institution, attributes of the individual proposer, geography, field of study); rigor of plans to address AI ethics (privacy, civil rights, fairness, transparency); approaches to minimize carbon impact of the work; originality and other relevant [criteria](#) considered for NSF merit review. In addition to criteria for gaining access, the NAIRR should develop criteria for removing users who do not adhere to an AI ethics code of conduct.

Regarding proposal solicitation NAIRR has two primary options: (1) issue and manage its own RFPs, in which case NAIRR will have more control over proposal selection, or (2) serve as a resource for existing research funding agencies to supply their awardees with compute and data. There are advantages to both options and Google does not have a strong preference; *however we do recommend that the NAIRR Task Force determine a best approach early on and articulate this so stakeholders can plan appropriately.*

D. Capabilities required to create and maintain a shared computing infrastructure

The NAIRR should recommend the Resource use open source and open-source-based technologies such as container abstraction layers, open APIs, public codebases, and open source databases. This would ensure operational and technical consistency across public clouds or private data centers and effective management of

infrastructure, applications, and data across the organization. Given that many participating researchers and institutions may be using a particular cloud provider for their research, this would ensure that any applications made available to users are compatible and readily portable in a multi cloud environment further enabling users to move projects across different cloud environments. For researchers that operate on premise, this approach would also allow them to leverage both cloud and on-premise technologies. This strategy enables users to take advantage of all the public cloud services and best-of-breed features according to their needs. Designing an open-source, interoperable platform is critical to achieve this and will also significantly lower the barrier for other users to replicate results (lack of reproducibility has become a bit of a crisis in some fields over the past decade). Google's early adoption of open-source technology has demonstrated that use of open-source and open-source based technologies lead to more innovation, more public benefit and more democratic use of technologies.

Further, the NAIRR implementation plan should adopt a multi-cloud, multi-ML framework-enabled strategy. Multi-cloud is an architectural approach that enables users to leverage the strengths, such as tools and platform services, of multiple providers for various purposes, and gives them freedom from a prescriptive architecture of one single cloud provider. This approach is made possible by using open-source [technologies](#) to manage containerized applications. Containers are the important layers for enabling different types of cloud-based software applications. Containerized applications can run on any cloud environment, and even on on-premise technical infrastructures. They separate software applications from the underlying hardware and operating system they run on, allowing them to be deployed in a modular and hardware-agnostic fashion. This would enable interoperability and several important benefits. First, it would facilitate participation of users with different configurations of legacy investments in on-premise technical infrastructure and existing cloud usage. Interoperability would allow users to continue to use their legacy investments to the fullest extent, while taking advantage of future NAIRR cloud-based infrastructure services from a variety of providers. Interoperability also enables more innovation. The less users have to worry about accounting for proprietary idiosyncrasies between the various cloud environments they want to deploy their application with, the more they can focus on their research.

Additionally, we encourage the NAIRR to take a **multi-ML framework approach**: i.e. to consider that there are many cloud AI services (hyperparameter tuning, explainability/interpretability, etc.) that do require more interaction with the ML framework than is typically provided by a container. Given that, the roadmap should reflect that the NAIRR should include enabling features to allow TensorFlow, PyTorch, JAX, and other frameworks to be used directly and require containers to adhere to a common specification that would support additional functionality with any ML framework, beyond just training or inference.

Finally, we believe the NAIRR should require that compute resources provided be, at a minimum, carbon neutral (all of the major US cloud contributors have committed to carbon neutrality in their operations; Google has gone a step further and committed to operating on 24/7 carbon-free energy by 2030).

E. Barrier to Accessing Government Data

The NAIRR should facilitate streamlined access for more researchers to both already public and otherwise inaccessible government data (through Data Commons) and to open source software (OSS). Even researchers who are not accessing NAIRR compute should be able to use the NAIRR for data and OSS access (with appropriate privacy protections). Publicly available, non-USG data and OSS should also be included on the platform, including data and open source tools from private sector, international, and state and local government sources, and the NAIRR should serve as an impetus to adopt a common approach to structuring and labeling data (where possible) so it is more useful for researchers (though unstructured data should also be included).

We recommend that the NAIRR co-locate an instance of [Data Commons](#) in all NAIRR clouds, which we would provide as an in-kind contribution. This would serve as the vehicle for making existing and newly available public data more useful and for providing a standards-driven data governance process. In order to effectively use the wealth of publicly available data (including data on data.gov) for AI (and other) research, a dataset needs to be processed – this involves locating the data, cleaning it, aligning the schemas of disparate data and ensuring [machine readability](#), etc. This expensive error prone process, which is repeated for each analysis, not only becomes a barrier to the use of data, but also leads to problems of reproducibility in research questions. Cleaning a large dataset is no small feat; before making Google datasets publicly available for the open-source community, we spend hundreds of hours standardizing data and validating quality.

Data Commons does the data processing once and makes the processed data widely available via standard schemas and Cloud APIs. Data Commons is not another repository of data sets (like data.gov or dataverse). Instead, it is a single unified database created by normalizing/aligning the schemas and entity references across these different datasets. So, for example, if a researcher wants the population, violent crime rate and unemployment rate of a county, the researcher does not have to go to three different datasets (Census, FBI and BLS), but can instead, get it from a single database, using one schema, one API. Co-locating updated versions of Data Commons with the NAIRR would therefore enable more effective use of the resource.

The data governance process should also consider whether the dataset is representative of the intended content. Even if data is usable and representative of some situations, it may not be appropriate for every application. Data made available to users should be accompanied by [data cards](#) where possible. In addition, NAIRR can ensure users have access to tools like [Facets](#) to analyze the makeup of a dataset and evaluate the best ways to put it to use. They should also have access to training on building more [representative datasets](#) as well as access to tools such as interfaces like the [Crowdsourcing application](#). If Google provided Data Commons as a data access platform, Google could also offer some of Google's AI Fairness tools as part of the platform. Similarly, if Google provided access to the [Vetex AI](#) platform, NAIRR users could leverage Google Cloud's MLOps tools (including model management), Explanations AI and associated fairness tools, and future model risk management, data and model governance and fairness tools.

The Task Force may also wish to consider using the NAIRR as a central hub for offering restricted access to some types of sensitive government data (for example, health-related data, financial services, and Census data), and to establish a process for granting such access. This could be modeled after the [Census Bureau process](#) governing restricted data access.

F. Security requirements & access controls

We recommend the NAIRR support high-level privacy principles and/or risk-based control frameworks that guide the inclusion of cloud service providers (CSPs). NAIRR should ensure CSPs adhere to international standards, such as the need to be certified against a minimum set of internationally-recognized standards that we anticipate will be available by the time this resource is made broadly available: ISO SC42 standards and the NIST AI Risk Management Framework. Voluntary Industry Codes of Conduct are also a helpful proxy for assessing the extent of a CSPs offerings in this area.

The success of a research initiative potentially involving sensitive data depends upon the ability to reliably credential users and provide granular access management. These challenges are further complicated in contexts that require that credentialing and access be managed across a range of research institutions, resource providers, and data sources. To address these complexities and the need to ensure a high level of data privacy and security, **NAIRR should consider zero trust principles and architecture** which provide greater security by requiring parties accessing data to demonstrate that they are who they say they are based on multiple, context-aware signals. **Embracing principles of least privilege**, in which parties accessing data are given access only to the resources that are needed to complete the task can similarly help balance NAIRR's grand research objectives with the need to maintain data and security. **Automated identity and access**

management systems can support the implementation of a strategy based upon both sets of principles at scale.

Finally, we recommend NAIRR consider including technologies that simplify the compliance configuration process and provide seamless platform compatibility between government and commercial cloud environments. The benefit of this is that it can quickly and easily create controlled environments where U.S. data location and personnel access controls are automatically enforced in any of our U.S. cloud regions. For example, Google's Assured Workloads can help NAIRR meet the high security and compliance standards through simple controls that help customers prevent misconfigurations and be more confident in compliance. Assured Workloads can be configured to meet a range of compliance requirements such as those set forth by the Department of Defense (i.e., IL4), the FBI's Criminal Justice Information Services Division (CJIS), and the Federal Risk and Authorization Management Program (FedRAMP).

G. Privacy & Civil Rights

Ideally, the NAIRR will give users access to new AI capabilities, including by expanding access to data such as government data that has been historically difficult to access but that would enable further research into key areas such as bias and fairness. This is because sharing certain types of government data and providing powerful AI tools raises important questions about how to protect people's privacy and rights. There are a number of steps the NAIRR can take to protect privacy and civil rights, including ensuring the appropriate expertise amongst its staff, reviewers, and users, and evaluating proposals for privacy and civil rights protections. These steps are outlined in more detail in our answer to question (3).

Public Cloud providers continue to develop innovative and **up-to-date privacy protections and emerging techniques to learn from sensitive data**. For example, Federated Learning is a technique for training global ML models without data ever leaving a person's device, which has been made available through open-source tools, such as TensorFlow Federated. Another technique is Differential Privacy, which can offer strong guarantees that training data details aren't inappropriately exposed in ML models. Additionally, researchers are experimenting more and more with using **small training datasets and zero-shot learning**.

Similarly, modern public cloud environments provide robust security by design, with native security capabilities that can help cyber defenders address weaknesses and capability gaps in existing security efforts. Too many legacy, on-prem systems continue to be expensive to maintain and hard to secure. Traditional security approaches of continually adding more tools, more people, more compliance have not worked. Furthermore, public

cloud providers can provide innovative solutions like confidential computing, which extends the protections provided by encryption to situations where data is in use.

H. Sustaining the Resource

We believe the NAIRR should be a multi-cloud hosting platform for commercial Cloud resources (as opposed to a new Cloud platform developed by government or academia). Building a new platform from the ground up would require a huge investment of dollars and expertise, and even once built would not have the advantages brought by the scale of existing Cloud providers (e.g., security, operational, and energy efficiency). A multi-cloud hosting platform model would allow the USG to negotiate rates and in-kind contributions from private sector partners. As outlined above, we believe that both academic and private sector users should have access to the same resources -- but not at the same cost. Rates should be lower and subsidized by the USG for academic and government users. Furthermore, as noted above, use should be prioritized based on the public benefit of the proposed research, including plans to publish the work in an open access format (either through an open access publication, journal, or version on arxiv or a similar platform). Any user who does not plan to publish the results of their work should not receive priority, and should be required to pay the full cost of the services.

In order to achieve significant impact, we recommend that the USG fund the resource at \$500 million/year or more. In addition, private sector participants could provide in-kind support for the program, for example through low-cost access for certain kinds of users (in particular academics and government researchers), training, and data. As mentioned above, Google would like to offer updated versions of Data Commons as an in-kind contribution. With program funding, the NAIRR should aim to dedicate a relatively small portion of the compute resources (e.g., 30%) to cutting-edge, large-scale ML research (i.e. requiring 1 exaflop+ of compute), and reserve the remaining resources for small- to medium-scale projects. We estimate that this structure would allow for a handful of large projects per year, and thousands of smaller-scale projects.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

NAIRR should prioritize facilitating public access to and the use of government data sets (including US federal, state, and local government data) through an updated version of [Data Commons](#) that is co-located with the NAIRR Cloud(s) to support the design, development, deployment, and operation of AI applications (more details under question 1(E).

In addition to data governance, the NAIRR should prioritize compute access for (a) a diverse group of users and (b) use cases that are likely to lead to public benefit. A merit review process for applications (described in more detail in section 1(c) should consider representation from federal, local and state governments; a diverse range of academic institutions, and startups. Public benefit should be measured by applicant plans to publish their work, by the likely contribution of the work to addressing social issues, and by the ethics considerations outlined in the proposal (see above).

Finally, NAIRR should provide access for users to existing, open access ML models, which users can fine tune for specific applications. This is important, because it will allow researchers to build on existing models and make more efficient use of compute resources.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

There are a number of measures that the NAIRR can adopt to reinforce AI ethics and responsibility. For example, the NAIRR should:

- Ensure ongoing engagement with government, civil society, and industry bodies working on AI ethics around the world to share and learn from experiences;
- Ensure in-house expertise spanning relevant areas: technical, ethics, human / civil rights, social scientists, Responsible AI researchers, cloud, legal, and public policy;
- Develop a code of conduct and training for all users, and ensure users demonstrate awareness and training in best practices for privacy and civil rights-related issues;
- Design review process(es) to evaluate individual use cases and build a knowledge base of best practices over time, with defined standing committee members pulling from in-house expertise above, ensuring technical and non-technical members that are multi-disciplinary and include social scientists, human rights experts, and tech ethicists;
- Appoint an advisory committee with representation from widely-recognized Responsible AI experts--both technical and non-technical--across academia, tech ethics, human rights, civil society and industry, and lead staff member to coordinate engagement processes on AI ethics issues;
- Develop and require AI ethics training for government-funded researchers (analogous to research ethics training required for bioscience researchers funded by the NIH);
- Require that proposals include a section on ethics, which can reference institutional ethics review processes, and encourage applicants to subject

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Hewlett Packard Enterprise

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Introduction

Hewlett Packard Enterprise (HPE) believes artificial intelligence (AI) can amplify human capabilities and that investments into the cyberinfrastructure that fuels AI research and development is in our best national interest. That's why we appreciate the opportunity to respond to the Request for Information (RFI) to assist the NAIRR Task Force in establishing and sustaining a NAIRR and propose a roadmap detailing how such a resource should be established and sustained. Driven by the promise of AI technologies, scientists across the country are eager to study AI and explore how it may benefit their work. NAIRR should capitalize on these scientists' enthusiasm and untapped potential by democratizing access to facilities, funds, data sets, knowledge, tools, and any other resources that would allow them to study and leverage AI. For this goal to be actionable, our responses target three key recommendations:

- **Prioritizing, enabling access to AI-ready datasets** – This includes providing clarity on where to find relevant data sets for use in priority areas such as climate change research as well as information about the data sets provenance and metadata to understand details of the data itself.
- **Leverage existing computing facilities and infrastructure wherever possible** – Advancing AI will also require significant computing resources. To maximize government return on investment in these resources, we recommend building on existing computing facilities and infrastructure whenever possible. Additional investments should ideally be targeted to augment these capabilities as needed. This balanced approach would accelerate the availability of NAIRR computing capabilities and also, importantly, enable the Nation to benefit from the decades of experience and expertise that have developed in our National Labs and other national computing centers.
- **Continue support for research into foundational AI methods** – AI is an active field of research, in which everything from processor design and model architectures to overall workload optimization and best practices for mitigating systemic bias remain open areas of innovation. Support for this foundational research should continue to be a priority, and NAIRR should seek to develop processes and pathways for testing and translating new innovations into the resources it makes available.

Response to RFI Questions Roadmap Elements

1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in) for which you are identifying options.]

Establishment Goals/Sustainment and Metrics for Success

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

The NAIRR resource has the potential to add value along the following dimensions for academic, industry and government stakeholders:

- Scientific – e.g., make discoveries possible
- Technological – e.g., create an ecosystem of innovation that includes for-profit and non-profit uses
- Data-access – e.g., democratize access to data resources (e.g. maps, weather, etc.)
- Scalability – e.g., ability to grow to more users, more partners, more contributors

- Collaborative – e.g., attract and develop talent

True success on the scientific, technological, access, scalability and collaboration dimensions can be measured using one or more of the following measures of success:

- Increased productivity and reduced time-to-insight
- Number of individual, small-business, non-profit, and corporate users
- Number of papers, press-releases, publications that acknowledge the resource
- Increasing activity among networks of vibrant communities of creative domain scientists and data scientists that can collaborate effectively
- Platforms to orchestrate/instantiate accelerated discovery workflows
- Increase use of AI in addressing high-priority, cross-cutting urgent national challenges such as climate change

Plan for Ownership and Administration

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

- i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and

Envisioning an effective and impactful national artificial intelligence research resource will require federal leadership at the top. Expanding the coordinating and convening role of OSTP and OMB, and to some extent, the NITRD around this resource is likely a good first step. Those agencies and efforts have been successful historically with coordinating federal efforts in areas such as high performance computing, large scale networking and cybersecurity. With Artificial intelligence R&D being a program component area for the NITRD for several years now, it makes sense to leverage ongoing activity and established intergovernmental relationships. Furthermore, the annual OSTP/OMB multi-agency research and development priorities memo continues to urge agencies across government to prioritize, promote and collaborate on AI. Keeping and empowering these mechanisms is likely to lead to faster progress than having to start up a new bureaucracy, from scratch. It also helps provide important top level direction and guidance in support of a national AI research resource.

That said, it is important to understand the evolution of AI research resources in the US. At HPE, we are intimately familiar with this evolution and see our company as a worldwide leader, particularly at the high end. We are experiencing firsthand how pervasive and rapidly evolving the field of AI is. In our experience, organizations across the country, at various levels of sophistication, are trying to use AI to improve their operations and activities. Most are overwhelmed with what they face. Keeping up with AI's rapid pace of change is exceedingly difficult for most academic, commercial, government and organizations. Knowing this should be fundamental when trying to offer a nationwide research resource.

- ii. A governance structure for the Research Resource, including oversight and decision making authorities

HPE recommends that an advisory board with membership in a diverse community of stakeholders be established to work with the office responsible for implementation, deployment, and administration.

To ensure the National AI research resource is broadly available throughout the United States, in urban, suburban and rural communities, and especially within communities that are traditionally underrepresented

in the development and use of technology, the federal government will need the participation of nearly every agency. AI will heavily impact all NITRD participating agencies and programs. But there are many more agencies and programs beyond NITRD that need to participate. The fastest way to accomplish this is to keep most of the responsibility for implementation, deployment and administration within existing federal programs and agencies with input from the advisory board to assure the full diversity of the community is recognized.

The big piece that is missing is the coordination piece. AI, including machine learning, deep learning, inferencing and other forms, is already computationally expensive and will get much more expensive in the coming years. Furthermore, data is also growing exponentially, in part because of inexpensive sensors and the ability to collect a lot of data and can be exploited in the use of AI. AI requires storing and moving large amounts of data. When looking at AI users across communities, we see enormous opportunities to share AI artifacts, models, data, computational results and more. This would produce dramatic savings in resources, including power and reduce potentially massive redundancies. A key role of a research resource would be to bring communities together to improve coordination between organizations at the many levels they operate.

Model for Governance and Oversight

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

No one new bureaucracy is likely to be capable of doing all this work in a reasonable amount of time. By leveraging existing agencies and their current oversight and decision-making operations, we can build on existing processes that have been honed over time, and augment with new organizational infrastructure only where needed. Some agencies, like the Departments of Energy, Defense and National Science Foundation can provide additional help with respect to the technical needs of a research resource. As such, HPE recommends taking a “whole of government” approach, similar to what was done under the National Strategic Computing Initiative (NSCI). NSCI had lead agencies (DOD/DOE/NSF) and deployment agencies (NOAA, NASA, etc.) all operating within the NSCI umbrella. NAIRR could build from this model.

Repositories for AI models, metadata, data and other artifacts will be an important component for a national research source. In that respect, we believe that an open-source community-based approach is better way to proceed. NAIRR could consider a “git-for AI data,” modeled after the wildly popular GitHub which provides open source software development and version control. The ubiquitous availability of git server in most IT environments across the United States eases the deployment complexity of AI resources since it relies on already available services and does not create an additional administration burden to shared infrastructure services. For example, git-like simplicity allows AI meta-data to be tracked and managed separately in a decentralized usage model, easily cloned in any environment, following an open-source community-based approach that enables wide adoption and reuse. From there, we can develop technology to build trustworthy AI artifacts that meet desired tradeoffs between accuracy and uncertainty.

Shared Computing Infrastructure

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

Creating and maintaining a shared computing infrastructure should leverage existing organizations that currently provide shared computing infrastructure to a wide community of students and researchers, including DoE computing user facilities and NSF computing centers. These centers provide many of the capabilities needed to support the NAIRR goals. However, because access to curated data sets is critically important but not all AI research needs large compute allocations, the NAIRR should supplement these with shared computing infrastructure that is designed for and operated with a view toward maximizing the number of projects that can use those datasets, and providing the educational tools and services needed to reach diverse communities of potential users, who may have no experience using existing shared computing infrastructure.

Features to emphasize:

- Open vs. proprietary technology in the resource
- Sufficient availability, scale, and capability of resources with a commitment to growth as needed
- Strategic approach that encourages market competitiveness and ongoing technological innovation without sacrificing inter-operability
- Seamless interfaces and processes for inter-platform multi-modal data exploration and interrogation
- Tools, queries, methods and algorithms to platforms and data

Assessment of Recommended Solutions

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

Over the past two decades, the federal government has made significant progress toward increasing the priority placed on data management and curation within federal agencies. These efforts were significantly bolstered by the 2017 Foundations for Evidence-Based Policymaking Act. Despite this progress and these efforts, however, accessing high-quality government data sets remains challenging. One of the key reasons for this is the large number of government portals and initiatives that have been launched, often changing with each administration. This lack of a persistent, global data infrastructure often leads to confusion over where to find data sets, which data sets reflect the most current data and preparation process, and what data are best-suited to various AI applications.

While the NAIRR will need to address questions of ensuring data lineage and provenance when making data sets available, more importantly it must address the challenge of establishing and maintaining a persistent solution for making data sets accessible and available to researchers. This will include establishing capabilities within federal agencies to manage data availability and curation that persist across administrations and through budget cycles. Critically, however, this global solution must also be developed in close partnership with the research community, ideally leveraging non-federally owned or managed infrastructure and capabilities in non-profit entities such as universities or other research institutes. At the same time, the NAIRR should be careful to ensure that data infrastructure does not create a preference for one type of technology or limit competition from future technology and/or data solution providers. To ensure this, the NAIRR will need to reconsider the role of the US Government in managing its data sets, emphasizing the government's role in convening stakeholders and driving toward shared best-practices while federating data infrastructure management.

Some of the challenges that NAIRR will need through address through broad stakeholder engagement include:

- **Data set discovery** – Making it easy to find data sets and guidance on which data sets are considered the “gold standard” in various disciplines
- **Sharing**– Who can access data and how?
- **Trust** – How can users know they are accessing the data they think they are accessing, understand their lineage, assess their “quality” (e.g. reproducibility etc.), and validate their provenance?
- **Security** – How can data sets be accessed and analyzed securely and in a way that maintains privacy?
- **Incentives** – Many critical data sets are not owned by the federal government. What incentives would help increase safe, trusted, and equitable data exchange?

Assessment of Security Requirements

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;

The identification, classification and protection of dataset categories, and the categories of the communities that are likely to access them, are the key aspects of NAIRR security requirements.

There are three classifications of assets:

- Strategic Datasets – dataset that provide national advantage and security, this includes critical datasets and models and computational techniques (and computational infrastructure and resources)
- Academic Datasets - datasets and models that are useful and beneficial to the academic community
- Community Datasets - learning datasets and models that are useful for training but not a basis for standard or trusted knowledge datasets

Similarly, there are parallel communities that create and consume these assets:

- Researchers who are authorized to create, consume and derive works from the national advantage and security assets
- Researchers who can leverage and derive works from the academic resources
- Teachers and learners who can consume and leverage training datasets and models.

The boundaries between the three categories are strict and the transmission of data from one to the other must be managed with respect to confidentiality and integrity.

Key security principles to consider are:

- Multifactor authentication to identify and categorize users and manage their appropriate access to the categories of assets
- Protection of the integrity of models and datasets and their inputs
- Encryption of data in motion and data at rest
- Recordation of state changes to assets with accountability to the level of a single individual.

Assessment of Privacy and Civil Rights and Civil Liberties

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

As it provides AI researchers and students across scientific fields with access to an advanced computing ecosystem and data infrastructure, the NAIRR has the opportunity and duty to expose researchers and students to holistic tools for assessing the ramifications of applying AI in their disciplines using legal and ethical frameworks which are authentic and inclusive to the populations they serve.

Across the US Federal Government and US-based industries and academic communities, a portfolio of tailored ethical frameworks has emerged that link organizational core concepts to AI ethics principles and then to operational commitments and specifications. The NAIRR should provide access to this expanding portfolio and guidance to researchers and students in selecting and tailoring their AI ethical and legal frameworks. Importantly, the NAIRR should seek to ensure the development of an AI culture in which ethical and legal considerations are foundational to every AI project rather than viewed as a “box-checking” compliance exercise.

Sustainment Plan through Federal Funding

H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; and

Sustained, incremental funding for NAIRR will be critical to its success. Securing this funding will take leadership from the White House, and in particular from Office of Management and Budget. But it will also be incumbent on individual agencies and programs to re-prioritize their funding proposals in a way that maximizes the use of new and existing research resources toward the broader goal of advancing our national AI capability in a way that continues to serve their core missions. For those agencies and programs already providing a significant amount of AI research resources for the benefit of government, industry and researchers, additional funding will likely still be needed to achieve the goals of the NAIRR. This is because the many programs that we are aware of are already highly oversubscribed and increasing demand will only further stymie progress.

Increasing funding for research resources incrementally will allow these agencies and programs to provide more resources to more users, building on their past experiences and best practices. Furthermore, the NAIRR should consider whether there are mechanisms for facilitating collaboration between grantees at different agencies that already fund AI or AI-enabled research. This kind of cross-disciplinary and cross-mission collaboration has the potential to rapidly advance progress in foundational AI technologies as well as in the way AI is used in a variety of applications by cross-pollinating ideas and enabling the sharing of best practices

Lastly, HPE is a strong proponent of federally funded partnerships involving the private sector. We see value in having the private sector involved in nearly all aspects of a national AI research resource ecosystem.

Establishment and Sustainment Parameters for NAIRR

I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

To the extent possible, HPE believes that the NAIRR should leverage processes, infrastructure, and best-practices from existing federal agencies and facilities for managing user facilities. Several science agencies, including DOE, NSF, DoC/NIST, and DoC/NOAA operate shared user facilities with competitively allocated access for use by American scientists. While the closest analog to NAIRR will likely be found in

the computing facilities operated by NSF and DOE, the NAIRR should broadly consider the practices by which federally-funded user facilities (including non-computing facilities such as photon sources and observatories) are made available to develop suitable operating approaches. Most importantly, the NAIRR should prioritize approaches to resource allocation that ensure broad-based access to all Americans, with an emphasis on ensuring that smaller projects and projects proposed by less experienced researchers are not de-prioritized.

To achieve this vision of “democratized” access, the NAIRR will likely need to be a federated resource across several agencies. Within each agency, the NAIRR may prioritize supporting infrastructure and resource access that best relates to that agency’s mission. A similar pattern currently exists with respect to high-performance computing resources, which are made available to the public by DOE and NSF with each agency targeting a different set of researchers and outcomes with the OSTP-coordinated Networking and Information Technology Research and Development (NITRD) providing loose coordination. NAIRR, however, should seek to provide a more robust coordinating function than NITRD currently allows, especially with respect to ensuring sufficient resources are made available in the budget across each agency.

Prioritized Capabilities and Services

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

To broadly advance the use of AI, the education of an AI-savvy workforce, and research in AI, HPE believes that the NAIRR should prioritize making high-quality curated datasets available, with access to expertise to enable use of those datasets. Data ethics and privacy best-practices should be emphasized in the data curation process as well as in the use of that data.

Additionally, HPE believes that the NAIRR should offer enough resource capacity at acceptable capability, that grows over time. Because requirements and technologies are diverse and rapidly evolving, a strategy that allows market competitiveness and the provision of diverse technologies should be adopted.

Other Priorities for consideration:

- Data ethics and privacy best-practices sharing (included in data curation)
- Coordinated access policy across all resources (including across agencies)
- Forum for convening (and also studying the interaction)

Principles of Ethical and Responsible Research and Development

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Under the joint leadership of Hewlett Packard Labs and the HPE Chief Compliance and Privacy Offices, a pan-HPE team has recently completed the drafting of the HPE global AI Ethics Principles. That team is now engaged in the even more challenging phase of operationalization of those principles into commitments and specifications to guide our team members, customer and partners in the utilization of AI across our products, our processes and our partnerships worldwide. While there is a large portfolio of AI Ethics principles and frameworks which have emerged from government, industry and academia, we felt that we needed to tailor a framework which was derived from our company purpose, to advance the way that

people live and work, and was authentic to our roles in creating, supplying and consuming AI technologies.

As we gain experience applying these principles in practice, we're uncovering gaps where conventional AI applied to real world situations cannot be applied with confidence in meeting our AI ethics principles, revealing instead issues of Bias, Explainability, Trust and Robustness. It's not just a matter of being more careful or deliberate with today's state of the art, these are technology gaps and closing them will take engineering and ingenuity, and this forms the basis of our Hewlett Packard Labs research agenda. That team is developing novel techniques and approaches to Model Synthesis and Analysis, the Data Foundation underpinning ethically robust AI, and Hardware Acceleration that will enable explainable, robust AI to be operated equitably and sustainably.

By embracing a dual mission of providing both providing access to a holistic advanced computing ecosystem and data infrastructure as well as to training and resources to enable researchers and students to establish AI Ethics frameworks tailored to their communities and concerns, the NAIRR could both enable those teams to both demonstrate with transparency how current technology can be applied with confidence and illuminate where current technologies fall short, and innovation is required. As it is intentionally focused on traditionally underserved communities, this dual mission would provide not only access to technology but add the lived experience of those communities directly into the ethical discussions concerning their use.

Without mandating a particular ethical framework, engaging with an NAIRR focused on the dual mission of technology and ethical rigor could be established as a benchmark for due diligence in ethical application of AI technologies.

Existing Building Blocks

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

A successful deployment of the NAIRR will require several foundational components which are already in existence. These are important building blocks that the NAIRR can leverage to its advantage. We recognize that as governments and industry actively consider how to approach governing AI, standards will play a key role in establishing the connection between rules and practical implementation.

HPE is currently engaged with a few government agencies that we believe are well positioned to provide the critical building blocks for the NAIRR implementation. These include the DOE National Labs, NSF advanced computing centers, NASA, NOAA and a few other agencies that have high end computing and data infrastructure as well as a surrounding ecosystem of system, applications, and other technical support. This existing infrastructure with some incremental funding could be leveraged to provide AI resources for a broader community. It is significantly cheaper to leverage an existing ecosystem, rather than starting up new centers from scratch.

There will be other considerations for the successful implementation that we'll identify for your consideration, as follows. Of critical importance are the data repositories. These will need to be federated, and to streamline access for users of the repository, may also need to be imported or securely linked into NAIRR resources. We believe some of the efforts by NIST to cultivate trustworthy AI systems, can also provide the complimentary foundational building block to the efforts the NAIRR Task Force with regards to secure data repository access.

We also recognize that there will be specific use cases where Public cloud providers will provide additional resources to supplement those directly engaged by the NAIRR. Often these can be leveraged for users just

getting started and may not have computing resources beyond the client devices. It is also important to build on existing arrangements with a proven track record like the NSF-funded CloudBank model, which provides individual researchers access to commercial clouds for NSF-funded research.

Finally, we strongly recommend engaging with Industry consortia like the ML Commons and other standards bodies such as Open Geospatial Consortium, as well as private sector solutions especially with highly data intensive applications like earth observations.

Public-Private Partnerships Roles

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

- Standards bodies
- Training/tutorials

HPE believes public-private partnership is not just a mechanism to enable capacity or access to experts. Rather, HPE has a long history of developing new technologies, in partnership with our customers, to enable new capabilities. HPE believes that the NAIRR researchers will identify gaps in capabilities that public-private partnership, through NRE funding, can fill. Therefore, there should be provisions allowed for the development of these capabilities through the auspices of NAIRR.

HPE and our government partners aspire to leverage current and future advances in artificial intelligence and high performance computing to dramatically accelerate scientific discovery and expand the frontiers of scientific knowledge to benefit US national security, scientific leadership, and industrial competitiveness.

Shared Principles:

- Keep AI-related software open source as much as possible and prevent vendor lock-in
- Establish an inviting, standards based open infrastructure environment to foster AI hardware and software innovations from the greater community. Standards proposed will take into account the need to ensure system-wide functionality with competitive performance relative to proprietary solutions
- Hide complexity from the user in order to make high end AI accessible to a wider number and range of domain scientists
- Greatly improve the collection, curation, storage, movement and analysis of AI relevant data at scale, taking into account the expected greater volume, velocity and variety.
- Enable real-time and near real-time control/steering of experiments and scientific based operations through the use of AI connecting the system of systems
- Emphasize the development of high productivity software and tools that perform well at scale

Lab Collaboration:

HPE believes that lab and other key customer collaborations will be critical for an AI for Science at scale initiative to succeed. HPE's efforts must be driven by advanced future use cases from customers in order to be relevant and attack the problems that matter most to our user community. Our large customers have extensive experience in advancing science and have access to ground breaking use cases, and connections to the researchers and the greater scientific community developing them.

Limitations to Democratize Access to AI R&D

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Efficiently executing AI workloads is a very computing and data intensive endeavor. In the last decade, the AI R&D community has aggressively moved to a specialized hardware and software infrastructure, primarily using GPU acceleration, to achieve the desired cost/performance target. Democratizing access to AI R&D requires an open approach to accessing this specialized infrastructure, and the tools (such as runtime software, programming environment, and libraries) that are required to operate it.

There are a few obstacles to making this happen that the NAIRR will have to consider. The AI accelerator market is experiencing a high degree of variability regarding openness in how accelerators are integrated into larger computing systems. In traditional accelerated systems, accelerators cards (such as GPUs, FPGAs, or AI-dedicated ASICs) are attached to standard servers through open interfaces (like PCIe) and multiple integrators can offer comparable systems. In newer, tightly integrated "AI appliances", entire systems are offered as proprietary black boxes and only accessible through network interfaces, thus limiting the end user choice and the integration points. A last, but not less important category of AI systems are those that are kept proprietary, not available in the open market, and only accessible through a specific interface. Among these are for example what is offered by cloud service providers with dedicated proprietary accelerators. In this case, the lack of openness raises to the software level, and while users are not impacted at the hardware integration level, they are constrained in the way in which they can access and orchestrate the accelerators, and where the data they process can live.

Overcoming these limitations requires a "community call to action" that the NAIRR could sponsor, to preserve a healthy and open ecosystem that offers AI R&D users choice at all levels. While proprietary innovation is fundamental for the AI ecosystem to thrive as a business, there are important dimensions that can benefit from an open approach, such as storage, **networking, virtualization, runtime, workflow, scaling and security**. For example, the NAIRR could sponsor a set of workshops or other activities with the objective of producing open recommendations (possibly standards) and a reference architecture, for future accelerated systems. By provide an open reference architecture to everyone as part of the requirement of future systems, the NAIRR could help create a blueprint for an open AI R&D infrastructure that can motivate all players to participate.


One last important challenge to the democratization of AI R&D is the lack of established approaches the combine AI and scientific computation. The main development of the AI frameworks (such as PyTorch or TensorFlow) is happening outside of the scientific community, driven by different players and incentives. The NAIRR could play a vital role to act as a catalyst for the scientific community to help define, design, and publicize desirable AI and HPC integration models, such as AI by the side, outside-the-loop, and inside-the-loop of HPC codes. Publishing open benchmarking practices and measurements for AI in science is another important set of objectives that can help AI R&D to establish common ground and be able to compare alternatives. Finally, working with the open-source community so that lower level of accelerator software can be penetrated by open-source compilers, optimizers, and programming tool chains.

Democratize access by allocating a portion of the resources to be made available through a light-weight request process, to enable researchers and students without a history of using large centralized HPC resources to start using NAIRR resources without substantial barriers.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Hyperion Research

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



Request for Information (RFI): Hyperion Research's response to RFI on an Implementation Plan for a National Artificial Intelligence Research Resource

Response Date: September 1, 2021

Organization: Hyperion Research

Point of Contact:

Mike Thorp
Sr. Global Account Executive
Hyperion Research

Contributing individuals:

Earl C. Joseph - CEO
Steve Conway - Senior Adviser, HPC Market Dynamics
Mark Nossokoff – Senior Analyst
Alex Norton - Principal Technology Analyst and Data Analysis Manager

Enclosed is Hyperion Research's response to the RFI regarding the development of an implementation roadmap for a shared research infrastructure that would provide Artificial Intelligence (AI) researchers and students across scientific disciplines with access to computational resources, high-quality data, educational tools, and user support.

Hyperion Research is a leading industry analyst, market research firm focused on the High Performance Computing (HPC) market and associated technologies, including deep insights and expertise relative to Artificial Intelligence (AI). We have current engagements and many past successful engagements with federal government organizations, including Department of Energy (DOE), Department of Defense (DOD), National Science Foundation (NSF) and various intelligence agencies. Hyperion Research hosts the HPC User Forum, since 2000, we have held 77 HPC User Forum meetings in the U.S. and around the world. The HPC User Forum was established to promote the health of the global HPC industry and address issues of common concern to users. The organization is directed by a volunteer Steering Committee of users from government, industry and academia, and operated for the users. (www.hpcuserforum.com)

We look forward to addressing any follow-up questions the NAIRR Task Force may have relative to our input. We would also ask to be considered for input for potential future requests and direct engagement with the Task Force.

Hyperion Research proposes the following recommendations to the RFI response questions:

RESPONSE SECTION:

This RFI seeks input from a broad array of stakeholders on the topics set forth below. Comments from the public will be used to inform the Task Force's consideration of options and development of an implementation roadmap.

Responders are invited to provide answers to the following questions (please number your responses accordingly):

1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

NAIRR Roadmap Element	Suggestions
A	<ul style="list-style-type: none"> • <i>Establish subcommittee to be led by a member of the Task Force to develop metrics for success. Commission a private study of NAIRR constituents to determine their needs and requirements for success.</i> • <i>To sustain a successful NAIRR roadmap and success, understanding what foreign entities are doing relative to AI is critical. Developing a formalized mechanism to identify and track non-US based technologies (HW, SW, tools, infrastructure) could provide the necessary global insights.</i> • <i>Once goals and metrics are established, developing a methodology for tracking NAIRR progress to attainment of those goals. This could also include comparison and progress against international organizations and countries, including China.</i>
B i.	
B ii.	The Covid-19 HPC Consortium is an excellent model off which to build a similar governance, oversight, and decision-making authority.
C	
D	<ul style="list-style-type: none"> • <i>Commission a private study to identify best practices and pitfalls to avoid for the creation and maintenance of existing shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country.</i>

	<ul style="list-style-type: none"> • <i>Create a consortium or advisory council of public and private sector researchers and organizations to share best practices and develop standardization and benchmarks of AI technologies. By bringing together multiple organizations with different approaches and skillsets, collaboration can facilitate deeper research and capabilities.</i>
E	<i>Tracking of progress of NAIRR goals. Tracking NAIRR progress compared to international organizations/countries, including China.</i>
F	<i>Commission a private study of NAIRR constituents to identify barriers and propose solutions for the dissemination and use of high-quality government data sets.</i>
G	
H	<i>Commission a private study of NAIRR constituents.</i>
I	<i>Commission a private study of NAIRR constituents.</i>

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

A non-exhaustive list of capabilities and services to be prioritized include:

- *Standardization and hosting of public data sets to enable broader access to high quality, large data sets for training.*
- *A concerted focus on developing the necessary AI skillsets in undergraduate and graduate-level courses and degrees. If there are more skilled researchers out there, there is less competition for building AI research teams.*
- *Access to the necessary hardware solutions, especially access to the latest and greatest accelerators from market staples to emergent technologies.*
- *NAIRR should heavily leverage and model the existing US infrastructure for advanced scientific and engineering research, especially the country's unrivaled network of NSF centers and DOE labs.*
 - *These centers/labs have long relied on HPC to support their advanced research. Hyperion Research studies, including for DOE, show that nearly all these sites already perform AI research to a greater or lesser extent. It therefore makes good sense in our opinion for the NAIRR to tie into this work.*
 - *America's next-generation exascale and pre-exascale supercomputers, intended for DOE and NSF sites, have been expressly designed to support both existing scientific-engineering workloads and new and emerging AI*

workloads. Hyperion Research studies funded by federal agencies or private-sector firms also support leveraging these current efforts.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Many corporations and organizations establish and communicate a set of corporate values to which they aspire and live. Following this model, the NAIRR could:

- *Identify a set of key "AI" values, provide examples, and actively, repeatedly communicate and promote them.*
- *Require partner organizations and constituents to also pledge to support and comply with the values.*

Additionally, the NAIRR could lead the effort to create a global set of standards and benchmarks by:

- *Developing a set of benchmarks to measure things like data bias for race and gender.*
- *Establishing a third-party organization to review and regulate outputs of models to ensure civil rights are not violated.*

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

There already exist some public data sets made available through published work and from government agencies, like the National Cancer Institute (NCI) and National Institute of Health (NIH). Continuing to expand those, as well as working out the funding and management of shared public datasets, is crucial to enabling future developments in AI.

Private companies like Google, OpenAI, IBM, Facebook, and others have created consortiums to impact standardization in certain areas of the AI space.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

a) Public-private partnerships should play a key role in the NAIRR. Collaboration between industry, academia, and government has been a hallmark of the HPC community for decades. One of the most recent and on-going exemplars is the

COVID-19 High Performance Computing Consortium. Under the leadership of OSTP, IBM, the Department of Energy, and other federal agencies, the consortium organized a broad array of public and private entities in an unbelievably short period of time to bring the appropriate resources to bear to combat COVID-19. In addition, providing access to the physical advanced computing resources, the consortium's efforts included governance, resource access, project approval process, and resource support.

- b) Review and leverage a prior US and global study performed on behalf of NCSA and NSF to identify and characterize public-private partnerships supporting advanced research. Utilize the findings of this study as a manual of best (and worse) practices for partnerships of this kind. A new, similar study could also be commissioned.*
6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?
- a. Private organizations can be protective and competitive when it comes to advances in AI capabilities. It will be important to establish the areas and rules of engagement for collaboration. These organizations need to understand the value of sharing best practices.*
 - b. Data privacy laws are getting stricter and stricter. Enabling shared data resources within the constraints of privacy laws is critical.*

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

IBM

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1, 2021

Wendy Wigen
National Coordination Office
Networking and Information Technology Research and Development
2415 Eisenhower Avenue
Alexandria, VA 22314

RE: IBM response, RFI on the National AI Research Resource [86 FR 39081]

Dear Ms. Wigen,

IBM appreciates the opportunity to comment on the Request for Information (RFI) on an [Implementation Plan for a National Artificial Intelligence Research Resource](#). IBM strongly supports its development.

IBM is an Artificial Intelligence (AI) and hybrid cloud technology leader and is engaged in research and development across a broad set of scientific and industry domains. IBM has extensive experience developing advanced computing for scientific research, including Summit, a 200-petaflop supercomputer built for Oak Ridge National Laboratory.¹ IBM co-created the COVID-19 High Performance Computing Consortium, where 43 members have carried out more than 100 projects.² IBM is also a leading provider of open source and open hybrid cloud architectures and technologies that simplify the integration of heterogeneous multi-cloud environments.³

If AI is to deliver on its full promise in advancing health, security and economic prosperity, democratization of its development by the research community and increasing the accessibility of both advanced computing and data will be key. Accordingly, the NAIRR must include the following core components:

1. Federated, hybrid cloud enabled computing resource – an accessible and easy-to-use hybrid- and multi-cloud computing resource built on open architecture that amalgamates various public clouds (such as Amazon, Azure), private clouds, and on-premises resources to create a single, unified, flexible compute infrastructure.⁴
2. Data and models – large scale, high-quality, trusted, AI-ready datasets and pre-trained AI models across the broad AI science and technology landscape.
3. Software and tools – integrated and interoperable software and platform technologies that support AI research and development and enable those with varying degrees of technology and science expertise to be productive.
4. Education – training materials, outreach activities, and user support that ensures easy, efficient, and effective use of the NAIRR.

If designed correctly, the NAIRR would be a pervasive, easily accessible federation of advanced computing resources, combined with a shared data infrastructure, that would bolster American leadership in AI research. Once again, IBM appreciates the opportunity to comment, and we look forward to future engagements. For any questions, please contact Mr. Jeffrey Brown, Science & Technology Policy Executive at jeffrey.brown@ibm.com.

Sincerely,

Dr. Dario Gil
Senior Vice President and Director of IBM Research

¹ “Summit Supercomputer Ranked Fastest Computer in the World.” *United States Department of Energy*, June 25, 2018. <https://www.energy.gov/articles/summit-supercomputer-ranked-fastest-computer-world>.

² “The COVID-19 High Performance Computing Consortium.” <https://covid19-hpc-consortium.org/who-we-are>.

³ “Hybrid Cloud Solutions.” *IBM Cloud*. <https://www.ibm.com/cloud/hybrid>.

⁴ “What is Hybrid Cloud?” *IBM*. <https://www.ibm.com/cloud/learn/hybrid-cloud>.



IBM Response to the Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource
White House Office of Science and Technology Policy (OSTP)
National Science Foundation (NSF)

1. What options should the Task Force consider for any of roadmap elements and why?

A. Goals for establishment and sustainment of a NAIRR and metrics for success;

To guide the development and sustainment of the NAIRR over an extended time horizon, IBM supports metrics that continuously track progress throughout the lifecycle of the computing resource, including its technical specification, development, testing, rollout, use, maintenance, and eventual upgrade. Similarly, metrics should be established to track the development, promulgation, and use of data resources across the complete lifecycle, including support for data provenance and versioning, and comprehensive management of metadata, classifications and policies. To overcome both compute *and* data constraints to democratize AI research and development, and therefore catalyze innovation, IBM proposes three categories of metrics:

1. **Researcher usage and productivity:** Initially, metrics should focus on the construction and rollout of the resource, including consumption of open hybrid cloud technologies and the availability of high-quality datasets and their use. As the NAIRR matures, metrics should shift to assess the overall impact of compute and data, including replicas of provided data sets, on researcher productivity. Specific metrics could include the number of institutions using the resource, the increase in scale of research (including amounts of data and compute employed), time taken to conduct and reproduce experiments, the number of research publications facilitated using NAIRR access, and shifts in the citation scores of researchers using the resource.
2. **Community development:** Modern scientific discovery demands the reproducibility of results, the creation of tools, standards or benchmarks, collaboration, and the effective communication of knowledge. Cultivating communities of discovery will be key to the NAIRR's adoption and spread. As the technical infrastructure is developed, parallel efforts should be undertaken to build a community of users whose actions and behaviors contribute to a virtuous cycle of increasing the mainstreaming and usage of the resource. Metrics could include increased participation in workshops related to the use of hybrid cloud for science or R&D, the development of tools for the community, the creation of new benchmarks adopted by the community, and the number of experiments shared within the community.⁵
3. **Wider impacts:** The NAIRR should result in long-term economic and job growth that bolsters the United States' global competitiveness. After the establishment of its infrastructure and community, metrics should be defined to assess breakthrough technologies and novel services unlocked by NAIRR. These metrics could include increases in productivity and automation, new products and startups created, and the integration of AI by large companies. Such economic benefits will inevitably have workforce impacts, and the resource should measure skills and jobs shifts spurred by its work. Finally, metrics should be developed to determine how the NAIRR has influenced international competitiveness.

⁵ A key element of many academic papers is comparison of the new approach to existing approaches, for instance, speed of accessing a data set. Having a standard yardstick to use for this comparison, which is used by many researchers, speeds the advancement of science. Thus, having the NAIRR foster such new benchmarks will contribute to speeding the advancement of science.



B. A plan for ownership and administration of the NAIRR, including:

i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource;

Deep, collaborative partnerships between government, industry, and universities have played a critical role in spurring innovation in the United States for decades. Federal investments in basic and translational research have enabled universities and federal partners to engage in high-risk, high-reward research activities, which have resulted in transformative ideas and commercialization that otherwise would not have happened.

While housing the NAIRR within a single government agency may streamline coordination and procurement negotiations, such a do-it-all approach could also impose short-term constraints regarding data sharing and impede the ability of the resource to grow flexibly over time. For example, a single agency approach could inhibit the pooling of data from a wider set of government agencies, limiting the inclusion of a broad range of diverse data sets as a critical ingredient for AI research and development.

To achieve success, the NAIRR needs to be endowed with flexibility that permits coordination with a diversity of stakeholders both inside and outside of government. And it needs to be built to sustain and evolve over the long run. IBM proposes a federated approach for implementation, deployment, and administration. A federated approach would allow sustained innovation and diversity in the constituent computing resources. For example, a federated approach would allow the NAIRR to take full advantage capabilities provided by the fast-advancing field of AI accelerator hardware and heterogeneous computing systems. A federated, virtualized approach could also better enable seamless access to computing and data infrastructure that allows researchers and scientists to participate in the procurement and deployment of the resource components themselves. Under an open hybrid- and multi-cloud model, the federal government should define how providers make participating computing and data resources available, for example, by documenting open hybrid cloud architectures, application programming interfaces, and standards for data and metadata representation.

ii. A governance structure for NAIRR, including oversight and decision-making authorities.

Given the complex nature of the task assigned to the NAIRR, IBM recommends that it be incubated as a Federally Funded Research and Development Center (FFRDC). The FFRDC model has been effective for similar roles.⁶ The NSF, for example, acts as an administrator for five FFRDCs, such as the [Science and Technology Policy Institute](#).⁷ FFRDCs further government research, technology development, systems acquisition, and policy guidance. The FFRDC model supports a “special long-term research or development need which cannot be met as effectively by existing in-house or contractor resources.” Furthermore, FFRDCs have “access, beyond that which is common to the normal contractual relationship, to government and supplier data, including sensitive and proprietary data, and to employees and installations equipment and real property.”⁸

An FFRDC governance structure should be sponsored and empowered by multiple agencies, including the NSF, the Department of Energy (DOE), the National Institutes of Health (NIH), as well as other scientific agencies, with strategic input from OSTP. Since FFRDCs marshal the expertise of government, industry, and academia to solve complex problems, the model is capable of meeting needs that cannot be met solely using government resources and contractors. Most importantly, the FFRDC model would allow multiple stakeholders (and government

⁶ “Federally- Funded Research and Development Centers (FFRDC’s): Background and Issues for Congress.” *Congressional Research Service*, April 3, 2020. <https://crsreports.congress.gov/product/pdf/R/R44629/6>.

⁷ “Master List of Federally Funded Research and Development Centers, by Agency and Type of Administration.” *National Science Foundation*. <https://www.nsf.gov/statistics/fedfunds/pubs/ffrdc/ffrdc.htm>.

⁸ *Ibid*.



agencies) to rally around its shared goal. And an FFRDC would ensure interoperability for communities of discovery to work collaboratively to solve common challenges.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

Ideally, a NAIRR FFRDC would define interoperability requirements including open interfaces for federation of the resource, and allow industry to compete in the development and fielding of leading-edge computing and data resources, as well as AI software, tools and platforms. Flexibility in procurement and use by researchers and “plug and play” by providers would require some degree of interoperability and standardization. However, given the wide diversity user needs and the fast-paced nature of computing and AI developments, care should be taken to avoid homogenization of the resource. A diversity of integrated computing resources would allow researchers to leverage resources across multiple clouds, utilize novel computing devices or hardware accelerators, and take advantage of supercomputers and high-performance clusters to meet their objectives.

To simplify procurement – which has been the Achilles’ heel of the FFRDC model – the NAIRR should learn from the NIH STRIDES program, which has greatly simplified cloud resource procurement for its researchers. But the FFRDC for the NAIRR should go further. The FFRDC should specify interoperable architectures and interfaces that allow diverse computing resources to be integrated into the federation. The FFRDC should employ open hybrid cloud mechanisms for managing the federation of multiple clouds (such as Amazon, Azure) and computing resources to allow seamless workflows and workloads to operate across the full diversity of provided resources. The FFRDC should also facilitate the integration of data management and AI software, tools, and platforms into the NAIRR to further aid and increase productivity of AI researchers.

To effectively manage data, users, and the provisioning of compute resources, governance of the compute and data resources should be automated where possible. To this end, it is important to bake in governance mechanisms that enable the automation of the management and the enforcement of policies, FAIR guidelines, regulations, and best practices. IBM suggests a “Software-defined Governance Framework” that can be easily configured and extended while ensuring consistency, transparency, and auditability through the lineage enabled via automation.

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

Data Sets/Secure Access Control: The NAIRR should adhere to the [FAIR guiding principles](#) for scientific data management and stewardship. FAIR provides guidelines to improve the findability, accessibility, interoperability, and reuse of digital assets. In line with FAIR, data made available through the NAIRR should be easy to find and read for both humans and computers. It should be machine readable to enable the automatic discovery of data sets and services.

Once a researcher has identified data sets, the NAIRR should provide direction on how the data can be accessed, so that the resource maintains overall control of security and authentication protocols. To the extent possible, enforcement of the security and governance requirements should be automated using available technologies.

Data acquired through the resource should be interoperable with applications or workflows for analysis, storage, and processing. Crucially, to allow for the reproducibility of experiments, the NAIRR should require interoperability so that workloads can run across diverse data and cloud environments. Further, data should be well-described so that it can be replaced or combined in different research experiments. And the NAIRR should require the replicability and portability of the data brought onto the resource. Replicability and portability need to be



enshrined to fend off “data gravity,” a phenomenon in which applications, computing, and users gravitate to a sole provider.⁹ Requiring interoperability, replicability and portability, in turn, ensures democratization of the resource and prevents overreliance on or lock-in by a single provider. The NAIRR will also need to provide tools and frameworks to enable moving and replicating data to enable computing over geographically distributed data sets.

Compute Resources/Scalability: To stand up and scale the NAIRR quickly, it should leverage commercially available compute resources. Commercial providers offer not only raw computing power, but sophisticated software stacks and user interfaces that have already been widely adopted by the AI research community. This would build on and broaden existing models such as NSF’s CloudBank.

Crucially, the NAIRR should adopt an open hybrid- and multi-cloud architecture, which is capable of scaling and delivering compute resources in a cost-effective manner. In short, the goal of hybrid- and multi-cloud is to establish a mix of public (e.g., Amazon and Azure) and private compute resources — and orchestration between them — that would give users flexibility to select the optimal resources for each application or workload and to move workloads freely as circumstances change. This may require making transitions between hybrid- and multi-cloud environments more seamless, for example, by employing a control plane that enables automation for optimal use of hybrid- and multi-cloud resources. Such interoperability would unlock new possibilities for researchers, including computational work that requires integration of data movement and computing across diverse locations.

But, just as the NAIRR leans on commercial providers to build the resource, it should also seek out inclusive frameworks to allow for the integration of existing public clouds. Multiple existing public clouds could be easily integrated into the NAIRR using an open hybrid cloud architecture. Advanced computing, accessible in hybrid- and multi-cloud usage models, will play a pivotal role in accelerating scientific discovery.

IBM has developed a hybrid- and multi-cloud approach for science that addresses key NAIRR requirements:

- *Heterogeneity* supports highly diverse resources including multiple public clouds (e.g., Amazon, Azure), on-premise infrastructure, and edge devices, which may include scientific instruments, sensors, physical devices, and entire labs and research organizations.
- *Consistency* enables portability, interoperability, and ease of management across heterogeneous environments.
- *Reproducibility* enables replication of scientific experiments and results regardless of differences in IT infrastructures or location of data and resources through the adoption of open standards and technologies.
- *Gravity* refers to the strong pull from extremely large data sets – some at a petabyte scale – as well as proximity required due to physical manifestation of experiments and instruments. IBM’s open hybrid cloud approach to data aims to optimize workflow deployment consistent with data gravity constraints and to avoid lock-in with specific storage environments.
- *Openness* refers to the prominence of open science practices that may dictate open – or *FAIR* – (findability, accessibility, interoperability, and reusability) data access, including encoding these principles into an open architecture for hybrid cloud through advanced technologies.

Such an approach, underpinned by advanced (and evolving) technologies, would make efficient use of taxpayer dollars, and would boost the usability of the resource. For example, adopting an open hybrid- and multi-cloud framework would allow a university researcher to connect their research experiment to the NAIRR. And government agencies could port their existing infrastructure to the NAIRR using the same hybrid cloud framework. This would make it easier for researchers using the NAIRR to run experiments that span cloud environments and multiple compute resources.

⁹ “What is Data Gravity?” *Data Centre Magazine*, November 2, 2020. <https://datacentremagazine.com/technology-and-ai/what-data-gravity>.



Educational Tools & Services/Resident Expertise: While some researchers are fluent in using tools such as hybrid cloud and AI to conduct research, many are not. Individual researchers, government entities, and cloud providers could benefit from engaging in technical exchanges, sharing best practices, and discussing challenges and opportunities of hybrid cloud for scientific research.

Educational tools and services and resident expertise must be built, including peer-to-peer knowledge sharing across different scientific communities to boost productivity across the board. To further this goal, the NAIRR should establish mechanisms for broad community sharing of best practices through annual conferences or other “birds of a feature” events. The NAIRR should also identify several important areas of shared interest across these communities and foster establishment of accessible testbeds that integrate computing resources, data, experiments, and evaluations for the specific AI application areas.

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

For data sets to be ready for AI, they need to be high quality and trusted, and they should be representative, transparent, and sufficient for training AI models employed in AI research. Additional challenges include privacy, dissemination, maintenance and automated governance and security policies. Given its ambition, the NAIRR faces numerous – but surmountable – challenges, including:

1. **Privacy, Protection, and Anonymization:** Amassing data from multiple government agencies to build and train AI models at scale faces overarching constraints imposed by the Privacy Act of 1974,¹⁰ which places limits on data disclosure and sharing. The Act granted certain data privacy waivers in the form of a research exception, which should be explicitly extended to cover the NAIRR and its work. To maintain compliance with the Act, government data gathered as part of the NAIRR should be stripped of personally identifiable information (PII). The NAIRR should consider utilizing mechanisms to automate the anonymization and removal of PII, and it should consider going one step further by ensuring that [differential privacy](#) is built into the resource. Furthermore, the NAIRR will need to apply security and compliance standards so that data can only be used for approved use cases. Enforcement of data privacy requirements should be automated to the extent possible using privacy-preserving computing techniques, such as [secure multiparty](#) and [fully homomorphic encryption](#).
2. **Intellectual Property (IP):** Datasets are often viewed as an IP asset, which could constrain data sharing between government, industry, and universities. IBM recommends that the NAIRR adopt protocols outlined in the Patent and Trademark Law Amendments Act of 1980, which would grant FFRDC partners IP rights to innovation produced from the NAIRR.¹¹
3. **Orchestration & Automation:** Given the volume of datasets that will be brought into the NAIRR, users will face the challenge of finding and using the right data for the right task. IBM proposes the adoption of a [hybrid data fabric](#) that takes a holistic view of the data lifecycle as it is created and used in federated distributed environments. A hybrid data fabric provides next-generation orchestration of secure and efficient data management. The hybrid data fabric is composed of a data plane and a control plane, which interacts with and manages data centric workloads in a brownfield, multi-vendor, multi-location environment.¹²
4. **Dissemination:** The NAIRR will need to not only curate data sets at an enormous scale and make them available via appropriate security and governance frameworks, it will also need to employ AI models – such as natural language processing – to solve problems in an open and community-driven way. For example, the [HuggingFace](#) Big Science effort uses large language models to solve problems in such a way.¹³ HuggingFace shows how a coordinated community effort supported by substantial computing resources (>5 million GPU

¹⁰ “Privacy Act of 1974.” *United States Department of Justice*. <https://www.justice.gov/opcl/privacy-act-1974>.

¹¹ “H.R.6933 – An act to amend the patent and trademark laws.” *Congress.gov*. <https://www.congress.gov/bill/96th-congress/house-bill/6933>.

¹² “Weaving data fabric into hybrid multicloud.” IBM Institute for Business Value. <https://www.ibm.com/thought-leadership/institute-business-value/report/data-fabric-multicloud>

¹³ “A one-year long research workshop on large multilingual models and datasets.” *Big Science Hugging Face*. <https://bigscience.huggingface.co>



hours) can take on difficult challenges like creating trusted open AI-ready data and large-scale training for language models.

5. **Data Scalability:** As the number of data sets on the platform grows and as data is combined, data becomes harder to discover and use. Therefore, the automated tagging and labeling of datasets with metadata can be important for their discoverability and findability by users. The NAIRR should consider supporting methods for automated metadata generation that analyze data sets and generate metadata tags that provide semantic information about the content of the data such as domain-specific terms and entities.¹⁴
6. **AI Model Trust and Transparency:** The NAIRR should incorporate mechanisms such as AI Fact Sheets that capture information throughout the data and AI model lifecycle that is critical for trust and transparency.¹⁵
7. **Self-Service User Access:** To reduce the time to value and enable scaling, researchers must be able to find the data they need. Building upon points 3 (orchestration) and 5 (scalability), a logically centralized catalog supporting semantic search and the enforcement of governance and access control will be needed. Such a tool will enable users to find specific data sets and related data sets, see the lineage and versions of the data they need, and make data sets visible only to authorized users.

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;

The NAIRR must be secure and resilient. Careful steps must be taken to ensure the confidentiality, integrity, and availability of both the computing resources and data provided by the NAIRR, and the computational workloads and derivative artifacts, such as intermediate data, results, models, and so on, from external attackers, administrators and providers of NAIRR, and other researchers. The number of constraints and requirements are too long to list in this RFI response (access control, encryption, authentication, logging, etc.), but we provide an outline below.

The NAIRR should be architected to be compliant with applicable standards commensurate with the classification of the data and workloads to be executed on it. For example, the base platform should ensure baseline compliance with NIST 800-53, provide the means of implementing the necessary controls, the ability to prove compliance, and allow auditors, developers, and users of NAIRR to verify compliance.¹⁶ Additionally, the platform must ensure compliance with additional regulations as necessary, such as HIPAA, FERPA, PCI DSS, and FedRAMP. Due to the open hybrid- and multi-cloud architecture of NAIRR, data and compute resources will be required to attest to their security compliance and integrity prior to admission, and workloads must be restricted to data and compute resources that are allowed by their classifications. For example, healthcare data cannot be processed on a system that is not HIPAA compliant. The resource needs to address unique security and privacy requirements, including:

1. **Multi-stakeholders and multiple administrative domains with a shared responsibility for security.**
2. **Multi-tenant:** Researchers represent different organizations and administration domains, and NAIRR must ensure isolation and separation of data and compute workloads.
3. **Encryption:** Data should be encrypted while at rest and in motion. The NAIRR will maintain full control over security keys and hardware security modules. Data may be encrypted using client-owned keys.
4. **Confidentiality of data:** Processed through secure enclaves and secure virtual machines, for example.
5. **Federated identity, federated authorization, and access management.** The platform should support the integration of policy-based data governance.

¹⁴ “IBM expands data and AI excellence with data catalogue technology in Cloud Pak for Data.” *IBM*, April 21, 2020. <https://www.ibm.com/blogs/journey-to-ai/2020/04/ibm-expands-data-and-ai-excellence-with-data-cataloging-technology-in-cloud-pak-for-data/>.

¹⁵ “How IBM is advancing AI governance to help clients build trust and transparency.” *IBM*, December 9, 2020.

<https://www.ibm.com/blogs/watson/2020/12/how-ibm-is-advancing-ai-governance-to-help-clients-build-trust-and-transparency/>.

¹⁶ “Security and Privacy Controls for Information Systems and Organizations.” *NIST Computer Security Resource Center*. September 2020. <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>.



6. Mixed access controls: a robust mechanism with a mix of access control models will allow for data sharing while maintaining security and privacy.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

The NAIRR should adopt [principles for trustworthy AI](#), including fairness, explainability and transparency, as they are developed by NIST as part of its [AI Risk Management Framework](#).

2. Which capabilities and services provided through the NAIRR should be prioritized?

- An open hybrid- and multi-cloud platform which can provide a consistent, seamless user experience across various public clouds (e.g., Amazon and Azure), private clouds, and on-premise resources.
- AI and data management software and workflow tools that simplify the use of cloud for science and R&D, including reducing the barrier to entry for researchers to adopt the cloud.
- Open standards and open technologies which accelerate interoperability, the adoption of cloud, and reproducibility for science and R&D use cases.
- Consistent methods and tools that automate data security and governance as well as identity and access management that can be used consistently across hybrid and multi-cloud platforms.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Building trust in the NAIRR is paramount. The ultimate beneficiaries of the resource – researchers and any member of the public benefiting from a research breakthrough – should be assured that findings have been reached in a manner that is ethical, responsible, and free of bias.

IBM has long adhered to [Principles for Trust and Transparency](#) to ensure that new technologies are developed and deployed in a transparent and explainable manner. And IBM has defined a risk-based AI governance policy framework, [Precision Regulation](#) that calls for fairness and security validated by testing for bias before AI is deployed and re-tested as appropriate throughout its use, especially in automated determinations and high-risk applications. It also suggests the designation of a lead AI ethics official, a model that the NAIRR should consider as it develops and scales a shared computing infrastructure.

To support [bias mitigation strategies](#), NAIRR should be proactive in creating and implementing AI ethics principles and practices, and ensure appropriate governance is in place to provide ongoing review and oversight of the research resource. Examples of tools to support bias mitigation and the trustworthy use of AI include the [AI Fairness 360 toolkit](#), [AI FactSheets](#), [IBM Watson OpenScale](#), and [IBM Watson capabilities designed to help businesses build trustworthy AI](#). Government, industry, and researchers will have shared responsibility to ensure that AI systems used as part of the research resource are tested and assessed for bias.

Further, the NAIRR must ensure that researchers from minority-serving institutions and those with limited research budgets are guide the science, design, and development and application of AI. While the AI field today may not reflect the demographics of our society, moving toward a more diverse AI research ecosystem could help to avoid and mitigate unwanted AI bias by including and reflecting the interests and values of diverse communities. The NAIRR should also support training that increases understanding, mitigation and recognition of bias and how it could be unintentionally introduced into AI systems during the development pipeline.



4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Academia and government have long pushed the boundaries of advanced computing — and have generated some of the most complex data and computationally intensive workloads along the way. At the same time, the private sector is leading the adoption of hybrid cloud, driven by software application modernization and digital transformation. To maximize the effectiveness of NAIRR, and to ultimately boost the productivity of the research community, it should combine building blocks from existing collaborative efforts to boost research cooperation, as well as leading-edge advances in the private sector.

First, public initiatives — such as the [European Open Science Cloud](#), NIH [STRIDES](#), and NSF [Cloudbank](#) — have used the cloud to connect researchers to datasets, software, and processing. Furthermore, public efforts have been made to cultivate communities of discovery that unite to create solutions to large-scale problems. For example, the [COVID-19 High Performance Computing Consortium](#) brings together academia, the federal government, and industry to provide access to some of the world’s most powerful high-performance computing resources in support of COVID-19 research. And the European Commission’s [Helix Neubla Science Cloud](#) has piloted a hybrid cloud platform for research, which feeds into the [European Open Science Cloud](#) that will create a shared research space for 1.8 million researchers.

These communities will capture the next dominant workflows and workloads for accelerated discovery and will drive a robust supply chain for innovation and value creation and achieve a scale of impact that is critical for society. Existing collaborative efforts should be boosted by private sector efforts such as IBM’s OpenShift, which allows for building and moving of workloads to boost responsiveness, scalability, and price.

In addition, there are many enabling technologies, most of which are based on open source, that help address some of the pain points that researchers have experienced under the above programs in their early stages of adoption of cloud. For example, hybrid- and multi-cloud platforms like OpenShift based on Kubernetes, and emerging technologies like Ray¹⁷ and IBM’s CodeFlare¹⁸ that further simplify the user experience for AI and data science on cloud can not only reduce the barrier to entry and enhance productivity for users but can also simplify the management and operations for providers. Emerging open technologies that help automate data security and governance, for example Fybrik¹⁹ or modular encryption, can be used to ensure that data is used and accessed according to policies defined by NAIRR.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

IBM recommends that the research resource be incubated using an FFRDC governance model. However, as the NAIRR evolves, it should consider leveraging public-private partnerships to accelerate its work. For example, in March 2020, OSTP initiated the public-private COVID-19 High-Performance Computing Consortium, a public-private partnership to marshal computing power to develop responses to the COVID-19 pandemic. In March 2021, IBM and the Cleveland Clinic announced a ten-year partnership to advance and apply open hybrid cloud architectures, AI, high-performance and quantum computing to accelerate health care and life sciences research.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

¹⁷ “Fast and Simple Distributed Computing.” RAY.io. <https://ray.io/>.

¹⁸ “CodeFlare drastically reduces time to set up, run, scale machine-learning tests.” IBM Research, July 7, 2021. <https://www.research.ibm.com/blog/codeflare-ml-experiments>.

¹⁹ “A native platform to control data usage.” Fybrik.io. <https://fybrik.io/>.



The NAIRR could face several roadblocks in its ability to democratize access to AI R&D:

Vendor lock-in: One of the greatest hurdles that researchers face as they begin to adopt the cloud is the difficulty of using more than one cloud platform for R&D. Researchers from NIH and other government agencies in particular point to the tremendous overhead of trying to operate across more than one cloud when it comes to secure data access and computing. We believe this problem can be solved through requiring interoperability through open standards and technologies that allow users to work more seamlessly across heterogeneous cloud environments to reduce this burden on the part of the user. For example, if data security, governance, and access management challenges could be made interoperable, regardless of cloud provider, it would allow researchers to become considerably more productive when moving from one platform to another, e.g. because of the availability of data in a particular location. In addition, over time we expect to see more and more development of algorithms that enable federated computation over geographically disjointed data sets. The NAIRR thus needs to support extensibility of the algorithms and frameworks it supports to break down barriers and fend off data gravity.

Inequity of funding: The high cost of compute resources and well-curated datasets means that researchers at large, well-funded universities benefit from access to these compute resources and well-curated datasets. However, smaller universities with finite research budgets often struggle to access – and make use of – these same resources. Lacking access to reasonably priced compute or data at scale will crimp the ability of the research resource to accelerate research and scientific discovery. While the NAIRR will theoretically increase access to compute and data, it must also be augmented by an increase in research and development spending that targets under-served communities. By moving to an open hybrid- and multi-cloud model that enables the integration of resources across multiple providers and on-premise IT, resources may become available at different price points but with the same user experience.

Human capital and user support: Boosting access to affordable compute and data is a major building block in democratizing access to AI R&D. But once the infrastructure of the NAIRR is in place, researchers will face the challenge of developing the skills and competencies needed to use the resource to accelerate their research. Closing the research gap with large, well-endowed institutions necessitates the creation of user resources and support communities that build human capital to operate and apply the capabilities of the research resource. To overcome these barriers to adoption and application, IBM proposes the creation of a research resource skills academy, along the lines of the [IBM Quantum Education and Research Initiative](#), which partners will 12 historically black colleges and universities to develop education, community resources, and technical communities to power the field of quantum computing.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

IEEE Standards Association

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Dear Joint Task Force Representatives,

The IEEE Standards Association (IEEE SA) acknowledges the White House OSTP/NSF Joint Task Force for its efforts to gather information for an Implementation Plan for a National AI Research Resource. We are pleased to provide input on the endeavor.

As background, the IEEE SA, is a globally recognized standards-setting body within IEEE. We develop consensus standards through an open process that engages industry and brings together a broad stakeholder community. IEEE standards set specifications and best practices based on current scientific and technological knowledge. IEEE SA has a portfolio of over 1,500 active standards and over 650 standards under development, including technical and impact standards relating to Artificial Intelligent Systems, next generation networks, IoT and Cybersecurity.

Thank you for the opportunity to contribute to your process. You may find our comments below.

Best regards,

Kristin Little
Senior Public Affairs Manager, IEEE SA

1. What options should the Task Force consider for any of the roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

Response
Regarding E, F, and G, IEEE SA posits that security and civil rights assessments are ethical assessments and therefore we recommend that the NAIRR conduct a profiling exercise to identify all ethical values likely to be at stake.
As a general comment, the elements identified provide a comprehensive overview of the issues to be addressed when implementing the roadmap. However, as a pre-condition, a comprehensive (regulatory) framework should be in place clearly prescribing NAIRR objectives in terms of ethicality. Such a framework could then be used to assess which data

should be used for what purposes and under what circumstances. The IEEE Ethics Certification Program for Autonomous and Intelligent Systems could be used for this purpose. This program is focused on criteria of Transparency, Accountability, Algorithmic Bias, Privacy and Governance. Each subject criteria set provides stratified criteria associated with adaptive-risk management properties that can be effectively applied to best understand the ethicality and human-centricity associated with AI systems.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Response
IEEE SA would suggest prioritizing the more strategy-oriented elements (C, E, F, and G) first, and then focusing on the more practical, implementation-oriented aspects (A, B, D, H, and I).
With respect to C and D, the requisite capabilities could be identified through the lens of an AI life cycle model. This would mean focusing on Design, Development, Deployment, Monitoring, and Decommissioning capabilities as appropriate.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Response
Once the capabilities are defined, a life cycle profiling of ethical aspects should be conducted and independently verified through an outside enterprise (IEEE could perform this function). Once verified, these values such as fairness, equity, bias, transparency etc. should be subject to an evaluation, monitoring, and regular verification regime.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Response
For your consideration, IEEE has various programs and activities that could serve as reference for building blocks for government, academic, or private-sector activities, resources, and services. Please reference: https://standards.ieee.org/initiatives/artificial-intelligence-systems/index.html

•

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Response
Partnership is essential to enhance credibility and ownership of the agreed approaches and ethical value preservation. At the same time, the use of public data by private parties and vice versa could pose problems if not embedded in a clear and ethical framework. Providing transparency, ensuring accountability, avoiding unfair bias, and offering the chance to opt out will be crucial.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?"

Response
Once the policies and processes for fair and equitable access are devised based on a creative profiling of the life cycle, the hazards to democratisation can simultaneously be noted and appropriate control measures identified. IEEE does not recommend having an <i>ex ante</i> list of such limitations since the impression would be that all barriers to democratisation are known and risks mitigated.

--
Kristin Little
Senior Manager, Public Affairs
IEEE Standards Association
445 Hoes Lane
Piscataway, NJ 08855 USA

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Indiana University

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



INDIANA UNIVERSITY

OFFICE OF THE VICE PRESIDENT
FOR RESEARCH

Email Subject line: RFI Response: National AI Research Resource

Email to: [REDACTED]

Final date of submission: October 1, 2021

To: Lynn Parker, Director of the National AI Initiative Office and Deputy U.S. Chief Technology Officer
Erwin Gianchandani, Senior Advisor for Translation, Innovation, and Partnerships, National Science Foundation

Subject: RFI Response: National AI Research Resource

Authoring organizational entities:

Indiana University Office of the VP for Research
Indiana University Office of VP for Information Technology
Pervasive Technology Institute
Center for Applied Cybersecurity Research
Indiana University Purdue University Institute of Integrative AI

Indiana University (IU) appreciates the opportunity to respond to the RFI on an Implementation Plan for a National Artificial Intelligence Research Resource, Federal Register vol. 86, no. 139, July 23, 2021.

Increasing the level of innovation in artificial intelligence in the United States will take a partnership between government agencies, public and private academic institutions, non-profit organizations, and private industry. We represent the research campuses of Indiana University (IU), a large public university system in the US Midwest. IU has two main research campuses: Indiana University Purdue University (IUPUI), an urban campus in the state capital, and Indiana University Bloomington (IUB), located in a small college town. IUB and IUPUI have a combined enrollment of over 70,000 students and include the Indiana University School of Medicine, the largest medical school in the United States. The two campuses have made sizable investments in technological infrastructure and digitized assets in support of research. This response represents our perspective on the unique capacity of an academic institution such as IU to contribute to AI innovation enabled by a National AI Research Resource and the national resource itself.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development for AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Respect for Civil Liberties: Respect for the civil liberties of especially vulnerable populations in the context of AI innovation is a complex technical, ethical, and philosophical problem. To make

meaningful progress in addressing these issues, they must be approached holistically, something academic institutions have the unique capacity to do. IUB and IUPUI contain and combine perspectives from the renowned Kinsey Institute for Research in Sex, Gender, and Reproduction, the Maurer Law School, the Luddy School of Informatics, Computing, and Engineering, the famed Ostrom Institute, the Center for Applied Cybersecurity Research, the Institute of Integrative AI, research centers focused on minority populations, and ethics programs working in health, business, IT, communications, and other fields.

Role of the CI Professional in NAIRR: CI Professionals enable the academic research carried out through the use of computational and data resources. As such, they are a key element of the overall integrity of a national AI research infrastructure. As AI models increasingly embody the same traits as the data that they use, the CI Professional must have tools and training to assess the inherent ethical limitations of one model over another, for instance. Models must have the ability to provide explainability when queried, and CI professionals be trained to interpret these results.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Academic Assets: Academic institutions such as IUB and IUPUI have made sizable investments in hardware and data infrastructure in support of research as well as operate significant infrastructure in support of their campuses overall. This technical infrastructure could be part of the NAIRR as a core building block by providing both computational infrastructure and real-world data. This combination of infrastructure, data, and capable CI staff could serve as a living laboratory to test and validate AI algorithms in a controlled environment with real-world data, particularly in focus areas that leverage the strengths of an academic institution. The strengths that IUB and IUPUI could bring to bear are in fundamental AI research and its application to the study of virus and human health as well as to global network management, cybersecurity, and to unique digitized assets.

Exemplar in NSF National CI Infrastructure Academic Service Providers: The National Science Foundation has funded a series of projects that provide national cyber infrastructure within the US. Through the Teragrid/XSEDE/ACCESS programs, what has arisen is a successful model of service providers, housed at academic institutions, that deploy and operate compute resources shared nationally, and who work together to provide common access mechanisms to federated compute resources. The service providers have additionally organized into a community of practice where they share knowledge. The agility and flexibility of this model was demonstrated by resources that were immediately available for COVID-19 research. IU has participated in these programs since the beginning and has rich expertise as a service provider.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Tools for hands-on AI education: To truly democratize AI such that we reach the diverse student populations in our public universities, AI education needs to improve by becoming more hands-on and more available. Students should be able to plug a new AI program into a robot, and see if

the robot dances better. Or plug a new AI algorithm into a vacuum cleaner to see if it cleans more efficiently. Or make an algorithmic change to a drone and experiment with collision avoidance. This improvement will require synthetic simulation environments available for classroom use. Large public universities with cyberinfrastructure investments such as IUB and IUPUI have the capacity to host these environments in partnership with commercial clouds.

Model commons: barriers to access can be lowered with investments in the creation of curated collections of data and AI models that are available using the principles of open science. This need is apparent to us for medical research to advance, though elsewhere as well. The NSF AI Institute (NSF 21-12606) effort in model commons could be a building block.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Infiltron

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

September 21, 2021

[REDACTED]

VIA ELECTRONIC SUBMISSION

Re: Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource

Infiltron is a Veteran-owned cybersecurity startup, focused on delivering proactive, real-time Internet-of-Things security solutions and improving biometric protection, fraudster detection, and malware intrusion for IoT devices. It leverages artificial intelligence (AI) technology, among others, to provide innovative cybersecurity protections.¹ Infiltron appreciates the opportunity to submit these comments to the National Artificial Intelligence Research Resource (NAIRR) Task Force.

Ethical, responsible, and fair development and deployment of AI technology (including democratizing access to AI research and development) needs to start at the top, with the government setting the right tone and parameters at multiple levels. The government has a vital role in creating legislation and regulation that holds companies accountable for responsibly developing AI technology.² As a customer, the government should expect its vendors to develop good technology rooted in good data. And—of particular relevance when it comes to the work of the NAIRR Task Force—the government can help startups access the quality data and resources they need to develop cutting-edge, responsible technology.

Infiltron seeks to hold itself accountable in developing responsible technology, for example through focusing on the accuracy of facial recognition data and AI trustworthiness. In addition to our responses to the questions posed, the NAIRR Task Force may be able to learn from Infiltron's experience, as noted below.

¹ Infiltron, <https://infiltron.net/>.

² For example, the Algorithmic Accountability Act considers how to combat bias in AI, protect the security of people's personal information, and monitor the accuracy and use of that data. *See* Press Release: Booker, Wyden, Clarke Introduce Bill Requiring Companies to Target Bias in Corporate Algorithms (Apr. 10, 2019), <https://www.booker.senate.gov/news/press/booker-wyden-clarke-introduce-bill-requiring-companies-to-target-bias-in-corporate-algorithms>. Infiltron has been involved in advancing similar state-level legislation in Georgia.

--

Response to Question 1. What options should the Task Force consider for any of roadmap elements A through I above, and why?

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

As noted above, startups need access to quality data and resources—the fundamental building blocks that will allow them to participate in developing cutting-edge, responsible AI technology. The NAIRR should be developed in a way that permits startups and other researchers access to multiple, unbiased data sets; that instills best practices across the innovation ecosystem; and that incorporates ethical norms during education and training.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

Diverse representation in the NAIRR’s governance will be critical to the development of fair, equitable, unbiased technologies. Diverse leaders should shape the NAIRR’s strategic direction, programmatic decisions, and allocation of resources. We know that diversity leads to better outcomes. Ensuring a broad set of stakeholders have a seat at the table will help the resource be dynamic and responsive to future developments and promote development of ethical, trustworthy AI. To this end, the Task Force should consider the importance of diverse representation along several lines—racial, ethnic, gender—as well as type of stakeholder—including startup founders or members of the startup ecosystem, academics, researchers, and others.³

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

Facial recognition technology has traditionally been developed by teams of primarily white and Asian men, and without representation from anyone of African descent or any Latino team members. And the data sets being used to train facial recognition AI are far from diverse enough. This lack of diversity has led to very real consequences of biased AI, from a man who was wrongfully arrested for a crime he did not commit to a young woman being improperly denied access to a skating rink earlier this year.⁴

³ See, e.g., #StartupsEverywhere: Chasity Wright, Founder and CTO, Infiltron, Engine (Nov. 11, 2020), <https://www.engine.is/news/startupseverywhere-warner-robins-ga-infiltron>.

⁴ See, e.g., Kashmir Hill, *Wrongfully Accused by an Algorithm*, N.Y.T. (June 24, 2020), <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>; Randy Wimbley & David Komer, *Black Teen Kicked Out of Skating Rink After Facial Recognition Camera Misidentified Her*, Fox2 Detroit (July 14,

This substantial potential for bias permeates across sectors where AI is already being used, like the Transportation Security Administration and U.S. Customs and Border Protection are using facial recognition technology at airports.⁵ The military is looking to use facial recognition for security at bases.⁶ To avoid bias in these and every other aspect of our lives, the government (along with everyone else) needs to be using the best software that is inclusive of all backgrounds, and with checks and balances on when and how facial recognition can even be used.

Infiltron works with diverse teams, and that diversity will show in what we are building—technology that attends to accuracy and defends against bias. And Infiltron is focused on biometric accountability,⁷ through our Biometric Legal Compliance and Governance Standard Protocol® and Biometric Technology Equality Standard®. Our software gathers and audits human biometric data, with an eye toward enhancing security, enforcing accountability, and mitigating risks (for example, legal and ethical risks).

The government must also prioritize diversity—in the NAIRR’s leadership (as noted above),⁸ as well as in the research teams it supports. But the problem, and the solution, run much deeper, to the data used to train AI systems. The NAIRR should prioritize providing multiple, unbiased data sets that ensure the broad array of all backgrounds are represented (for example, data sets that can look at hues of skin color). Startups and AI researchers need to be able to leverage those sorts of data sets when training AI systems.

As a customer and end-user of technology, the government should also expect its vendors to provide responsible technology that was developed through the use of good data. For the NAIRR’s purposes, it could consider building in stipulations about how this government support should (and should not) be used. And the Task Force should maintain open and active lines of communication across the government—from policymakers to procurement—so that lessons learned through the resource about responsible development and deployment of AI can benefit others, and likewise so that the NAIRR’s resources and best practices can be informed by other branches of government with relevant experience.

2021), <https://www.fox2detroit.com/news/teen-kicked-out-of-skating-rink-after-facial-recognition-camera-misidentified-her>; Dave Gershgorn, *Black Teen Barred from Skating Rink by Inaccurate Facial Recognition*, The Verge (July 15, 2021), <https://www.theverge.com/2021/7/15/22578801/black-teen-skating-rink-inaccurate-facial-recognition>.

⁵ David Oliver, *Facial Recognition Scanners Are Already at Some U.S. Airports*, USA Today (Aug. 16, 2019), <https://www.usatoday.com/story/travel/airline-news/2019/08/16/biometric-airport-screening-facial-recognition-every-thing-you-need-know/1998749001/>.

⁶ Singh Bisht, *U.S. Army Calls for Facial Recognition Tech to Secure Bases*, The Defense Post (Apr. 6, 2021), <https://www.thedefensepost.com/2021/04/06/us-army-facial-recognition-bases/>.

⁷ *Biometric Accountability*, Infiltron, <https://infiltron.net/biometricaccountability/>.

⁸ *Supra* response to question I.C.

Response to Question 2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

AI trustworthiness assessments are critical, and should be included among the NAIRR’s priorities. AI trustworthiness refers to evaluating AI systems against their stated solution, looking at things like “accuracy, explainability and interpretability, privacy, reliability, robustness, safety, and security (resilience) and mitigation of harmful bias.”⁹ For example, autonomous vehicles need to recognize stop signs even if they are marked with graffiti or obstructed by a tree branch. The car still needs to know that there is a stop sign there, and respond accordingly.

As the NAIRR Task Force considers what infrastructure the government will put in place, it should consider working towards better trustworthiness assessments. For example, the government can develop best practices or tools small companies can use—when they are just launching or getting started—to assess the trustworthiness of their AI solutions. Companies, including startups, need ways to test their AI, to make sure it is dependable and that it does not open up gaps for, e.g., hackers.

Response to Question 3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

The NAIRR can play a critical role promoting ethical, accountable, and trustworthy AI, and presents a unique opportunity to instill best practices that will last throughout product and company life cycles. As noted above, including diverse individuals in the development and governance of the NAIRR; ensuring multiple, unbiased datasets; and providing tools and best practices for trustworthiness can help promote these goals.¹⁰

Response to Question 6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Most innovative companies have (and will continue to have) some component of AI or leverage AI in their products. Ensuring startups access to the NAIRR would help more small companies compete and find success by lowering barriers to AI development. At the same time, with the proliferation of AI in business and as noted in above, the NAIRR can play a critical role in promoting best-practices for developing AI that is trustworthy and unbiased.¹¹ The NAIRR Task

⁹ *Overview: Artificial Intelligence*, NIST, <https://www.nist.gov/artificial-intelligence>.

¹⁰ *Supra* response to question 1.

¹¹ *Supra* responses to questions 1, 3.

Force should facilitate broad access for startups and other researchers, in order to achieve its aims of democratizing access and ethical AI development. To that end, it should also conduct proactive outreach to startups and members of the startup ecosystem, including with a special focus on historically underrepresented founders, to further the impact of the NAIRR.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Information Technology Industry Council

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

September 30, 2021

Ms. Wendy Wigen
NCO
2415 Eisenhower Avenue
Alexandria, VA 22314

Via email

Re: ITI Response to Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource (NAIRR)

The Information Technology Industry Council (ITI) appreciates the opportunity to respond to the White House Office of Science and Technology Policy's Request for Information on an Implementation Plan for the National Artificial Intelligence Research Resource.

ITI represents the world's leading information and communications technology (ICT) companies. We promote innovation worldwide, serving as the ICT industry's premier advocate and thought leader in the United States and around the globe. ITI's membership comprises leading innovative companies from all corners of the technology sector, including hardware, software, digital services, semiconductor, network equipment, and other internet and technology-enabled companies that rely on ICT to evolve their businesses. Artificial Intelligence (AI) is a priority technology area for many of our members, who develop and use AI systems to improve technology, facilitate business, and solve problems big and small. ITI and its member companies believe that effective government approaches to AI clear barriers to innovation, provide predictable and sustainable environments for business, protect public safety, and build public trust in the technology.

We recognize that AI is an active area of research that is constantly evolving and improving. To harness this growth, it is vital to both utilize AI's potential benefits while monitoring its impacts carefully, and research and development (R&D) is a critical enabler of those possibilities. Indeed, our *Global AI Policy Recommendations*, released earlier this year, include an entire section devoted to facilitating innovation and investment in AI that emphasizes the critical role R&D must play in that effort. As such, we welcome the opportunity to provide input on the Implementation Plan that the NAIRR Task Force is developing to guide the NAIRR.

We offer our response to several specific questions set forth in the RFI below.

Responses to Specific Questions

1. *What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]*

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

Metrics will play an important role in establishing and sustaining the NAIRR. We propose three distinct sets of metrics that we believe will be helpful in measuring success:

Researcher usage. We believe that focusing first on the development and rollout of the NAIRR will be most useful. Measuring the use of hybrid-cloud technologies, as well as consumption of high-value datasets, will help provide insight into how much the NAIRR is being utilized. Once the NAIRR has had the opportunity to develop, metrics could then shift to the impact of the resource on researcher productivity. For example, it may be useful to measure the number of institutions utilizing the resource, the increase in the scale of research (including the amounts of data and compute used), the time it takes to conduct and reproduce experiments, the number of research publications facilitated using NAIRR access, and shifts in the citation scores of researchers using the resource.

Community development. As the technical infrastructure for the resource is developed, the Task Force should simultaneously seek to develop a community of users whose consistent utilization make the NAIRR a mainstream, go-to resource. Metrics could include the development of tools for the community, the creation of new benchmarks or standards adopted by the community, and the number of experiments shared within the community.

Wider impacts. The Task Force should also seek to measure the wider impacts of the NAIRR, as it seeks to spur research by increasing access to compute power and high-quality data. Once the NAIRR has been established, it would be helpful to assess new technologies or services that have been developed as a result of using the resource. These metrics could include increases in productivity and automation, new products and startups created, and the integration of AI by large companies. This will lead to improvements in the U.S. economy, including to the workforce. The Task Force could use metrics like number of jobs created or shifted or skills acquired to deepen understanding of the wider impact. Metrics could also be developed to determine how the NAIRR has influenced international competitiveness.

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

ii. We encourage the Task Force to take a federated, or shared, infrastructure approach to implement, deploy, and administer the NAIRR. This will allow for a diversity of compute resources while enabling continued innovation. Indeed, such an approach will enable the NAIRR to rapidly upgrade/integrate new

capabilities, while also enabling researchers to participate in the procurement and deployment of the resource. We encourage the federal government, in standing up the NAIRR, to develop a standard set of interfaces for federation of the resource – this could include how providers make computing and data resources available, for example, by specifying hybrid cloud computing architectures, application programming interfaces, and standards for data and metadata representation.

Initially, we recommend that the NAIRR be governed as a Federally Funded Research and Development Center (FFRDC), which is a private sector entity contracted by the government to undertake research and sponsored by a specific government agency. As FFRDCs bring together stakeholders from government, industry, and academia, they are often able to provide perspective that other, more traditional government resources cannot. Because the NAIRR will be a complex, first-of-its-kind resource, we believe this will be a valuable governance structure. Importantly, the NAIRR should not only be sponsored by one agency, as is typical of most FFRDCS, but should instead be sponsored by multiple agencies, so as to ensure partnership with a breadth of entities that can both contribute to and benefit from the NAIRR. Beyond that, allowing multiple government agencies to sponsor the FFRDC will help drive coordination and facilitate data-sharing between agencies, which is a necessary prerequisite to standing up the NAIRR.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

As referenced above, we believe the FFRDC model will be most appropriate for governing the NAIRR. However, to the extent possible, we encourage the Task Force to consider how to automate day-to-day governance of the NAIRR. Indeed, this will be essential to ensuring researchers can easily access and utilize the NAIRR, helping to democratize access. One way this may be possible is to utilize a software-defined governance framework, which could be configured/reconfigured as necessary.

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

Data access/secure access control. Given the immense amounts of data that will be collated within the NAIRR, capabilities around data access are necessary. When devising guidelines for the NAIRR, the Task Force should keep in mind the FAIR Principles for data management, ensuring that data is findable, accessible, interoperable, and reusable. Indeed, in order to appropriately leverage and democratize the NAIRR, data needs to be both identified and consumed by humans. As such, the Task Force should ensure that there are directions set forth for how data housed in the NAIRR can be accessed. Because data brought into the NAIRR will

be diverse, it is also important that the Task Force consider how to make the data interoperable and also how to ensure replicability and portability. Without this, it will be difficult, if not impossible, for researchers to utilize the NAIRR.

Compute resources/scalability. Compute resources will also bring important capabilities to the NAIRR. Beyond housing data, the Task Force should consider how to facilitate access to computing resources. We recommend taking an approach based on shared infrastructure, in which the NAIRR leverages existing cloud infrastructure to enable more rapid deployment. Starting or building the computing resources from scratch would not be cost effective or efficient. A hybrid or multi-cloud approach would also allow for public cloud to be integrated into the NAIRR.

Educational tools. Educational tools are also necessary for ensuring the success of the NAIRR and democratizing access. The Task Force needs to develop educational materials that demonstrate how to utilize the NAIRR, including how AI models might be developed using the NAIRR in a responsible and ethical fashion. The Task Force should consider developing educational materials applicable to all levels of students, researchers, etc. to ensure the broadest reach. The Task Force should also consider how to set up mechanisms to share best practices around use of the NAIRR, bringing together users to share their experiences to facilitate broader uptake. This could take the form of an annual conference, or some other sort of symposium.

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource; and F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;

Data is fundamental to innovation in AI. Indeed, one of our key recommendations in our *Global AI Policy Recommendations* is for policymakers to consider how to increase access to government sources of publicly available data in machine-readable formats and across borders to enable access to a foundational building block of AI. While the federal government is already making vast quantities of data available through resources like data.gov, and both NOAA and NASA provide public access to petabytes of earth and space data, the United States does not have a national strategy on data access and data-sharing for public interest applications. It is not always clear who owns data or how much data belongs in the public space. Beyond that, because of inconsistent data management practices, data set size, compute costs, bandwidth limitations, and inconsistent use licenses, government data can be difficult and costly to use or combine with other datasets.

Additionally, there are challenges with making available large data sets that contain personally identifiable information. Although there are constructs like HIPAA which protect health data, similar controls do not exist for data collected in other settings, like faces and license plates captured by street-view mapping processes or video doorbell monitors. There may also be security challenges associated with the NAIRR, as a consolidated data resource may be a target

for malicious actors seeking to access and use the data for nefarious purposes. These challenges limit the innovation economy and can hinder academic research.

In order to address these challenges, we encourage the NAIRR Task Force to consider how to facilitate the adoption of AI by encouraging data sharing with meaningful stakeholders and determining how to best make datasets available to the broader AI research community. One way in which the NAIRR could do this is by taking a “data-as-a-service” (DaaS) approach, which sets forth clear policies and practices related to data access, security, retention, and use, with specified processes for updating those policies over time as technologies and data evolve. Although this approach has historically been used to monetize datasets, using this model for research could help to enable access to more data for AI modeling and research, while ensuring that safeguards are already built in. This model allows data to be located anywhere and searched from anywhere, while simultaneously reflecting an access structure that is based on graduated security policies that reflect the risks associated with each dataset. Data is “fluid” but controlled. The NAIRR Task Force should define overarching policies for dataset inclusion in the NAIRR, with dataset owners maintaining control of cost, access, and retention of their data within the scope of those overarching policies.

Additionally, the NAIRR should create opportunities to collect and distribute data responsibly, allowing U.S. citizens to opt-in to data collection by providing meaningful consent through services they already use, such as: healthcare exchanges, tax payments, social security, and Medicare. Another way to achieve this goal is to broker more data-sharing agreements, similar to the UK Open Banking Initiative or the IBM Mastercard Partnership.

We also encourage the NAIRR Task Force to develop flexible guidelines that are based on security and access control standards and best practices. Indeed, flexibility will be key in ensuring that such guidelines remain up to date as both cyber and privacy risks evolve. We believe the NAIRR should take a risk-based approach to security and access control, where data that poses minimal risk is made widely available and easily accessible. In doing so, the NAIRR should also consider developing guidelines that will help to protect data, intellectual property, and computing resources commensurate with the level of risk associated with those resources.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

We are encouraged that the NAIRR is considering privacy and civil liberties requirements in the context of responsible R&D, as we believe that maintaining appropriate privacy protections is imperative to fostering trust in AI technology.

Like all technologies, AI operates in an existing policy and regulatory framework, and accordingly, personal data and related privacy concerns must be taken into consideration. We face important questions around striking the right balance between various objectives in the responsible development of AI, such as ensuring accountability, which requires some level of visibility into an AI system, while also protecting privacy. Adding to the complexity of a dynamic international environment of laws that are not always in alignment, the US currently has a patchwork of privacy policies and regulations that could become more complex and

fragmented as additional states follow California’s lead in establishing state-level comprehensive consumer privacy laws (already, Virginia and Colorado have followed suit, and many other state legislatures are currently considering such bills). Conflicts across these laws could have a chilling effect on AI advancement, as well as other data-driven technologies. To maximize the use of AI, we need strong, globally accepted privacy standards to enable trust and interoperability, and to incentivize investment in research to develop new techniques for even stronger privacy and security guarantees. To achieve this, we recommend the development of a national privacy law in the United States, consistent with *ITI’s Framework to Advance Interoperable Rules on Privacy*.¹

In lieu of such a law, however, we encourage the NAIRR to play a leading role in evaluating and defining research data practices, AI use, and implementation guidelines that protect individuals’ privacy and ensure equity. One way in which the NAIRR could do this is to create a working group or advisory committee that includes academic and government researchers, computer scientists, industry representatives, non-profit organizations such as the Center for Democracy and Technology, and legal and ethics scholars with an understanding of modern computing technologies. Such a group of diverse, expert stakeholders could help the NAIRR to identify the privacy, security, civil rights, and civil liberties risks associated with the aggregation of vast amounts of data and the application and advancement of AI, and develop recommendations to address those risks. It could also lead and engage in public discussion around the risks that may be created by data and AI. Upon developing an understanding of what the specific risks might be, the NAIRR should seek to implement the recommendations of the working group, leveraging a similar combination of subject matter experts and stakeholders. Finally, the group should seek to develop and publish educational tools and guidelines for broad public understanding of the benefits of and risks inherent to data and advancing AI.

H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; and

The government plays an important role as an essential source of funding for long-term, high-risk research initiatives, and we recommend investment in diverse fields of AI research including cyber-defense, data analytics, fraud detection, robotics, human augmentation, natural language processing, and visualization and perception technology. Advanced algorithms, specialized computing hardware, high-quality data, and, most importantly, skilled human expertise are essential to enabling machine learning and the success of AI. To remain the leader in AI R&D, the United States must continue to promote an entrepreneurial environment, research network, and openness to talent. One of the reasons the United States has succeeded in this space is that it has invested heavily in R&D. In 2019, for example, the United States led the world in total R&D expenditures, with combined public and private sector spending totaling \$657 billion.² Maintaining American leadership in AI related R&D efforts will not only require continued government R&D investment, but in promoting scientific

¹ https://www.itic.org/public-policy/FINALFrameworktoAdvanceInteroperableRules%28FAIR%29onPrivacyFinal_NoWatermark.pdf

² https://stats.oecd.org/Index.aspx?DataSetCode=MSTI_PUB

collaboration among like-minded nations. Additionally, maintaining leadership in AI R&D will require continued strong political support from Congress and the Executive Branch, as well as active participation from the private sector and society.

The NAIRR should be primarily funded using public sources, as the NAIRR is intended to be a public resource. We encourage the U.S. government, and the NAIRR in particular, to map out a clear framework including funding commitments and timelines. One promising example of a framework the Administration could look to as a model, and develop in partnership with industry, is Canada's Pan-Canadian Artificial Intelligence Strategy, which is delivered through the Canadian Institute for Advanced Research (CIFAR). CIFAR received \$125 million to launch the Strategy in 2017, and the Canadian government has approved \$349 million (USD) to fund the Strategy over the next ten years.³ Through the Digital Europe and Horizon Europe programs, the European Commission plans to invest €1 billion per year in AI. It plans to seek additional investments from the private sector and Member States to reach an annual investment volume of €20 billion over the course of ten years.⁴ If the U.S. is to remain the global market leader in AI, it needs a strong government framework to integrate resources and set goals to enable AI to grow and prosper. Partnerships with the private sector have an important role to play as well – we encourage the NAIRR to look to such partnerships to supplement public funding, either in cash or in kind, for specific projects.

As a further means to supplement public funding, the NAIRR could consider offering access to data as a subscription-based service and access to compute resources based on a tiered pricing model. It could also consider rewarding researcher contributions to open resources (e.g., algorithms, research data sets) and providing technical consultants to optimize projects for computation. Importantly, though, the NAIRR should offer fee waivers or otherwise consider a structure that would take into account institution size, need, etc. so as not to inadvertently hinder access.

1. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

We believe that NAIRR should prioritize developing and maintaining a shared computing infrastructure, which requires compute resources. In order to do so, we encourage the NAIRR to leverage existing governmental, academic, and commercial cloud computing and data resources sourced from multiple vendors, which will enable the NAIRR to offer the latest computing resources to users. The NAIRR should develop a common abstraction layer that enables users to develop AI systems in the same way across all vendors and seeks to adopt open standards for both data and compute resources, allowing for seamless transition between vendors.

³ <https://vectorinstitute.ai/2021/05/03/federal-government-renews-pan-canadian-ai-strategy/>

⁴ <https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence>

Education is also a key capability deserving the NAIRR’s focus, given it is critical to the ethical and responsible deployment and use of AI. Developing cross-disciplinary educational materials that demonstrate how AI models can be developed using the NAIRR in a responsible and ethical manner would be a helpful first step. Such materials should be accessible to students at all levels (ranging from primary to graduate students and researchers) as well as other researchers.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Government, and the NAIRR in particular, has an important role to play in advancing ethical and responsible R&D of AI. In our *Global AI Policy Recommendations*, we emphasize the importance of facilitating public trust and understanding of AI technology, and we believe that the NAIRR can help to do so. We believe that the NAIRR can play an important role in facilitating interactions between AI companies and the communities they impact, with a view to better aligning stakeholders. More specifically, the NAIRR should invest further in research that supports the responsible development of AI, including in areas that improve the accountability, safety, fairness, and privacy of AI systems.

Beyond that, as the Task Force considers the funding and budget for the NAIRR R&D, we recommend that the NAIRR devote funds specifically to translational work – or the work it takes to translate basic research into industry-relevant insights or applications -- as a separate budget item. Too often in industry-academic engagements, it is assumed that research speaks for itself, which can lead to overburdened researchers doing significant amounts of work to ensure that industrial partners benefit from research, or only focusing on applied research where more basic, fundamental research is called for.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

There are important building blocks that the NAIRR can leverage. For example, as governments and industry actively consider how to approach governing AI, standards play a key role in forming a bridge between written rules and practical implementation. We believe that the NAIRR should seek to support industry engagement in rules-based, consensus-based international standards organizations that are developing AI standards, and should work with industry stakeholders to consider how to leverage those standards.

Governments should maintain technology neutral policies that limit mandatory implementation requirements (e.g., for public safety considerations) in favor of voluntary implementation and self-attestation. To the extent compliance requirements are established, they should adhere to international best practices of conformity assessment. ITI and the InterNational Committee for Information Technology Standards (INCITS) recently launched an initiative, “Standards as a Tool

for Achieving Public Policy and Regulatory Goals (SPUR),”⁵ and developed a list⁶ of commonly used and developing international standards for AI for U.S. policymaker to reference so they don’t have to start from scratch. For example, ISO/IEC 24028 and ISO/IEC 24027 provide information related to trustworthiness in AI systems including transparency, explainability and bias, which are great tools considering the many policy debates around AI. There are also various developing standards around the use of biometrics, focusing on facial recognition and fairness, such as ISO/IEC 22116 and 19795.

NIST has also been undertaking significant work to cultivate trustworthy AI systems, which we believe is a fundamental building block for the efforts the NAIRR Task Force will undertake. Indeed, the benchmarks, metrics, and standards that NIST is developing – including its AI Risk Management Framework -- should inform specific guidelines the NAIRR Task Force develops to govern the administration of the NAIRR itself.⁷ As such, we believe the NAIRR should seek to heavily leverage this work.

In response to question 2, we also highlighted the importance of leveraging existing cloud infrastructure to develop and maintain a shared infrastructure. We reiterate that much of this infrastructure already exists – there are high-performance computing clusters on several university campuses (there are HPCs at Yale University, Kettering University, Clark University, Columbia University, as a few examples), the private sector already has cloud computing infrastructure, and organizations like Folding@home aggregate compute resources.⁸ The Task Force should consider how it can leverage this existing infrastructure to administer the NAIRR. In particular, the NSF-funded CloudBank model, which provides individual researchers access to commercial clouds for NSF-funded research, could be especially relevant and applicable here.⁹

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public-private partnerships should play a critical role in the NAIRR. The National AI Advisory Committee and Subcommittee on AI and Law Enforcement, as well as the many public comment opportunities that have been provided to stakeholders as the US government seeks to implement the *National AI Initiative Act of 2020*, have been positive steps toward forging better collaboration between industry, government, and academia on AI. Given the rapid development and adoption of AI technologies in the commercial space, the need for consistent dialogue between the government and the private sector to inform research priorities, from both technical and social impact perspectives, cannot be understated. We recommend a regular cadence of dialogues such as quarterly discussions between the public and private sector. We encourage the NAIRR Task Force to continue to identify ways to incorporate regular

⁵ <https://www.incits.org/contentAsset/raw-data/7a702e19-e075-400c-92cb-88ad83fda0d1/reportFile/fb9b572d-49f2-4b4b-8074-a84f3edad942.pdf>

⁶ <https://www.incits.org/contentAsset/raw-data/688802a6-4aeb-4333-9782-0c8b855ba040/reportFile/bd433a99-972d-4af0-8c2f-cc712a82516d.pdf>

⁷ <https://www.nist.gov/artificial-intelligence>

⁸ <https://foldingathome.org/?lng=en-US>

⁹ <https://www.cloudbank.org/>

private sector participation to ensure AI is advanced in a fashion that is broadly beneficial to all Americans.

Many emerging AI technologies are designed to perform a specific task, such as assisting human employees or making tasks easier. Our ability to adapt to rapid technological change is critical and we must continue to be prepared to address the implications of AI on the existing and future workforce. By leveraging public-private partnerships – especially between industry partners, academic institutions, and governments – we can expedite AI R&D, democratize access, prioritize diversity and inclusion, and prepare our workforce for the jobs of the future. We recommend that government tap into the commercial space and when appropriate, form more public-private partnerships to maximize the potential of AI. For example, the Japanese government set up its country’s biggest research center through the government-backed Riken Institute, which involved 20 companies and research entities, with the goal of developing applicable AI in the medical and financial fields within 10 years.¹⁰ Another example of a successful public-private arrangement is the COVID-19 High Performance Computing Consortium, which brings together public, private, and academic participants to facilitate access to high-performance computing power to conduct COVID-19 research.¹¹ Similar consortia may be helpful to entertain as follow-ons or complements to the initial FFRDC-governed resource.

As the NAIRR matures, we encourage the NAIRR Task Force to look to the Information and Communications Technology Supply Chain Risk Management (ICT SCRMM) Task Force as a potential exemplar. The ICT SCRMM Task Force has been a highly successful public-private mechanism, which has developed products and tools to address some of the most pressing supply chain security challenges in the United States. As such, we encourage the Task Force to look to the structure of the ICT SCRMM Task Force as a guiding example of how a public-private partnership on AI may be similarly arranged, with private sector and government co-chairs leading the group and considering what working groups might be necessary to address challenges that may emerge as the NAIRR grows.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Facilitating Access to Datasets. As we mentioned in response to question 1E, we believe the success of many promising uses of AI will depend to a large extent on the availability of training data. However, the lack of a comprehensive data-sharing strategy in the United States may make it more difficult for the NAIRR to democratize access to AI R&D. Therefore, we recommend that the USG, with assistance from the NAIRR, develop a balanced framework for the responsible use of data. By leveraging large and diverse datasets and increased computing power and ingenuity, AI developers and other stakeholders are empowered to innovate across industries to find solutions that will meet the needs of individuals and society in unprecedented ways. AI driven medical diagnostics can alert doctors to early warning signs to treat patients more capably. Increasingly intelligent systems are capable of monitoring large volumes of

¹⁰ <https://asia.nikkei.com/Business/Technology/Researchers-to-develop-Japanese-style-AI>

¹¹ <https://covid19-hpc-consortium.org/>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Internet2

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



September 29, 2021

The National Science Foundation
Attn: Wendy Wigen, NCO
2415 Eisenhower Avenue
Alexandria, VA 22314
[REDACTED]

Re: RFI Response – National AI Research Resource (86 FR 39081)

Dear Ms. Wigen:

The University Corporation for Advanced Internet Development (d/b/a “Internet2”) is pleased to submit these comments in response to the National Science Foundation’s (“NSF”) request, in coordination with the Office of Science and Technology Policy (“OSTP”), for community input to inform the work of the National Artificial Intelligence Research Resource (“NAIRR”) Task Force.

BACKGROUND ON INTERNET2

Internet2 is a non-profit, member-driven advanced technology community founded in 1996 by the nation’s leading higher education institutions that provides a secure high-speed network, cloud solutions, research support, and identity and access management services tailored for research and education (“R&E”). Internet2 helps U.S. R&E organizations to solve shared technology challenges and develop innovative solutions in support of their educational, research, and community service missions. Internet2 also operates the nation’s largest and fastest coast-to-coast national research and education network (“NREN”), which now serves 323 U.S. universities, 59 government agencies, and 45 regional and state education networks. In addition, Internet2 collaborates with numerous leading corporations that work with the R&E community, as well as a multitude of NREN partners across the globe that represent more than 100 countries.

In addition, Internet2 operates the InCommon Federation, which facilitates secure, unified, and seamless single sign-on access to local and global research and academic collaboration resources for more than 10 million users and 800 educational institutions, research organizations, and commercial resource providers in the U.S.

InCommon makes possible trustworthy academic collaboration that reaches far beyond what a single organization can do on its own, through identity and access management technologies and services that are integrated across the globe.

Internet2 also offers eduroam service in the U.S to enable seamless roaming Wi-Fi at nearly 1,000 colleges, universities, schools, museums, libraries, and research facilities across the country.

Further, Internet2 plays a key role as a convener and facilitator of the R&E community. Internet2 regularly brings together representatives from academia, federal agencies, and private industry to foster collaboration and find solutions to common challenges related to R&E cyberinfrastructure that no single institution or organization could accomplish independently.

Finally, Internet2 supports the R&E community through the Internet2 NET+ Cloud Services Program, which enables R&E institutions in adopting cloud solutions through a streamlined process that minimizes the business, legal, financial, and other challenges associated with migrating from on-campus to cloud-based solutions.

INTERNET2'S RESPONSES TO SPECIFIC RFI QUESTIONS

1. What options should the Task Force consider for any of the roadmap elements A through I and why?

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.

To facilitate access to the envisioned advanced computing resources for researchers across the country, several key characteristics should be included in the establishment of the NAIRR:

- The NAIRR should support hybrid models for advanced computing resources to enable the broadest possible set technology advancements, scalability, and use cases ranging from batch processing to low latency real-time applications. The hybrid models should include on-premise, national centers, commercial clouds, in-network, and edge computing scenarios.
- Given that artificial intelligence (“AI”) and machine learning (“ML”) often require the use of large data sets, high-speed networking and data transfer

capabilities will be critical for all researchers across the country no matter their institution. Therefore, the NAIRR should support future federal efforts focused on funding such network capabilities across all higher education institutions.

- To enable global trustworthy academic AI collaborations, secure access should be based on the well-established model of Multilateral Federated Identity Management.
- To develop resident expertise, the NAIRR should include support for a variety of training programs and communities of practice for researchers, campus research computing and data professionals, and campus identity management professionals. Training the campus professionals is necessary to ensure a sustainable ecosystem of resident expertise for researchers.

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource.

Given hybrid models for the NAIRR, a fundamental barrier to the dissemination and use of high-quality government data sets is the lack of high-speed networking in many of the country's higher education institutions, especially those that have been chronically underserved. Consideration should be given to funding the necessary network infrastructure to connect all of these institutions. Additionally, commercial clouds are likely to be an important part of the NAIRR for both compute and storage resources. The associated data egress fees have been seen as a barrier by the research community. National and regional R&E networks have a long-proven track record of providing the networking infrastructure for data intensive research (including data egress waivers and private/direct connections for commercial cloud usage). Therefore, it is recommended that consideration be given to funding for the expansion and/or upgrading of R&E networks to ensure that all higher education institutions have adequate connectivity to the NAIRR resources and data sets.

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls.

The InCommon Federation (<https://incommon.org/>) implements security requirements for access controls based on several pertinent international standards for R&E federations. These should be considered for NAIRR access controls:

- [REFEDS MFA Profile](#), a means for RP/SP to request MFA and for the CSP/IdP to respond.

- [REFEDS Assurance Framework](#), especially for its identity proofing and credential issuance claims.
- [REFEDS R&S Entity Category](#), a program for release of basic user attributes in a data minimizing manner.
- [Security Incident Response Trust Framework for Federated Identity \(Sirtfi\)](#), a framework for collaborating in managing security incidents across R&E federations.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Secure Access Control:

The previously mentioned InCommon Federation is a well-established building block to support the NAIRR's needs for *secure access control*. InCommon facilitates secure, unified, and seamless single sign-on access to local and global research and academic collaboration resources for more than 10 million users and 800 educational institutions, research organizations, and commercial resource providers in the U.S.

InCommon makes possible trustworthy academic collaboration that reaches far beyond what a single organization can do on its own, through identity and access management technologies and services that are integrated across the globe.

Dissemination and use of high-quality government data sets:

To provide necessary high-speed network access to higher education and other research institutions and support the dissemination of high-quality government data sets, the NAIRR should leverage the ecosystem of state, regional, and national R&E networks. As the U.S. NREN, Internet2 provides private direct connections to commercial cloud providers and data egress waivers that can be used in support of the NAIRR. Additionally, Internet2 has collaborated with the NSF-funded Open Science Grid ("OSG") to host caching servers in the Internet2 national network, which, if expanded, could be used to disseminate high-quality government data sets.

Resident Expertise:

Internet2 provides various opportunities for training and project execution for identity management professionals at universities and research institutions, including [InCommon Academy](#) and the [InCommon Catalysts](#) program.

To address challenges and facilitate the most effective use of cloud technologies and platforms, Internet2 also offers the [Cloud Learning and Skills Sessions \(CLASS\)](#) program. CLASS assists U.S. higher education institutions, research IT organizations, and research groups by providing training to effectively leverage cloud technologies and platforms for research workflows. A combination of vendor-neutral guidance across cloud providers, and training on the tools and technologies supported by public cloud providers, allows a broad range of research use cases to utilize these important resources more effectively. Participants join the CLASS community of practice where they can share information and lessons learned.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public-private partnerships should not only play a role in the implementation of the NAIRR compute and storage resources, but they should play a role across all the capabilities required to implement the NAIRR.

Not-for-profit companies, such as R&E networks, have a long history of successfully partnering with government, academia, and the for-profit private sector to provide research cyberinfrastructure.

Examples can be found in NSF programs supporting cyberinfrastructure. The NSF Campus Cyberinfrastructure (CC*) program solicitation (NSF 21-528) Program-wide Criteria states:

The plan should include the campus status and plans with respect to *federated identity and specifically InCommon*, including: if the campus is registered with InCommon as supporting the Research and Scholarship (R&S) Entity Category to streamline integration with research applications.

NSF 21-528 also states under the Campus Computing and the Computing Continuum Program Area High-Performance Network Connectivity and Specification:

Proposals must describe the network connectivity of the proposed computing resource, both intra-campus [for example, the campus network path(s) connecting

the resource with the researchers and driving science applications on campus], and inter-campus (for example, showing the network path *connecting with the regional exchange point or Internet2*).

Furthermore, the NSF International Research and Education Network Connections (IRNC) program solicitation (NSF 20-535) states under the Infrastructure Improvement and Support Program Area – Open Exchange Points that:

Proposals must demonstrate a commitment to operation of an open exchange point, for example: support for homing of multiple international links; high capacity connectivity to Internet2, ESNNet, and, if relevant, FABRIC; providing maximum flexibility in connectivity and peering; automated, dynamic switching network services; and, in the best interests of the end users – the researchers, educators, and students in the U.S. – a demonstrated commitment to a productive partnership and collaboration with *Internet2, the primary NREN for the NSF community*.

Another example of public-private partnerships between Internet2, government, academia, and the for-profit private sector is the Cooperative Agreement (Award #1904444) between NSF and Internet2 called “Exploring Clouds for Acceleration of Science (E-CAS).” The Abstract states the following:

Internet2 leads the “Exploring Clouds for Acceleration of Science (E-CAS)” project in partnership with representative commercial cloud providers to accelerate scientific discoveries. The effort seeks to demonstrate the effectiveness of commercial cloud platforms and services in supporting applications that are critical to growing academic and research computing and computational science communities, and seeks to illustrate the viability of these services as an option for leading-edge research across a broad scope of science. The project helps researchers understand the potential benefit of larger-scale commercial platforms for scientific application workflows such as those currently using NSF’s High-Performance Computing (HPC). It also explores how scientific workflows can innovatively leverage advancements provided by commercial cloud providers. The project aims to accelerate scientific discovery through integration and optimization of commercial cloud service advancements; identify gaps between cloud provider capabilities and their potential for enhancing academic research; and provide initial steps in documenting emerging tools and leading deployment practices to share with the community.

Additionally, a recent example of a partnership between Internet2 and government is the NSF Cyberinfrastructure Center of Excellence Demonstration Pilot (Award #2137123) awarded to

Internet2 to support the Minority Serving Cyberinfrastructure Consortium (“MS-CC”). This Center of Excellence (“CoE”) Demonstration Pilot will initially support the advancement of research cyberinfrastructure capabilities and support systems for Historically Black Colleges and Universities (“HBCUs”) and Tribal Colleges and Universities (“TCUs”), with the goal of expanding to Hispanic Serving Institutions (“HSIs”) and other Minority Serving Institutions (“MSIs”). Internet2 will serve in a facilitating role by coordinating efforts that will allow participating organizations to work together to address common challenges that may be more difficult for each to resolve independently. The abstract states the following:

A key outcome of this grant is the formalization of a vibrant community of practice across MS-CC campuses that involves collaboration on cyberinfrastructure, education, and research applications. This grant enables HBCUs, TCUs, HSIs and other MSIs to accomplish together what they cannot do separately. The MS-CC is broadening participation in science, technology, engineering, and mathematics (STEM) by historically underrepresented groups in the United States’ research enterprise, enabling new perspectives to emerge and expand capabilities for the nation. MS-CC is advancing our nation’s economic growth, national security and global prosperity in ways that reflect the unique expertise and talent from HBCUs, TCUs, HSIs and other MSIs.

There also are examples of public-private partnerships between Internet2, regional R&E networks, and the National Institutes of Health (“NIH”) related to high-speed networking for the dissemination of NIH data, secure access controls for NIH’s Electronic Research Administration Portal, and NIH’s Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability ([STRIDES](#)) Initiative, which “allows NIH to explore the use of cloud environments to streamline NIH data use by partnering with commercial providers.”

Internet2, together with Mid-Atlantic Crossroads (MAX), the regional R&E network led by the University of Maryland, provides research network capacity to NIHnet (NIH’s backbone network). As described by NIH (<https://www.cit.nih.gov/services/networks>):

NIHnet is a high-speed, highly available network that interconnects NIH, the commodity internet, and the *Internet2 research network*. NIHnet runs a 100G backbone – which runs up to 50,000 times faster than the Internet access for most U.S. households – that provides fast, secure, consistent connections for network traffic, even when there are hardware or software disruptions or slowdowns.

Internet2 also is working with NIH to implement its upcoming requirements for multi-factor authentication (“MFA”) and identity proofing levels in connection with federated access to many of NIH’s services. NIH’s Electronic Research Administration (“eRA”) Portal began requiring

users to sign in with MFA on September 15, 2021. This change affects everyone who accesses eRA. Internet2 prepared an eRA Readiness Guide to IdP operators in the InCommon Federation to understand how this change impacts campus users. It also helps campuses prepare support options to assist users during the transition (<https://spaces.at.internet2.edu/display/federation/get-nih-ready>).

In support of NIH's Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability ("[STRIDES](#)") Initiative, Internet2 partnered with NIH and Google Cloud Platform ("GCP") to allow NIH-funded researchers to leverage Internet2's NET+ GCP terms to use a single agreement to access GCP for their enterprise and research needs as part of the initiative. Internet's NET+ program provides custom terms of service for cloud services developed by Internet2 higher education members specifically to meet the unique needs of research and higher education institutions.

All of these cited examples demonstrate a long track record of successful public-private partnerships between regional and national R&E networks, government, academia, and the for-profit private sector that can be used as models for many of the NAIRR capabilities outlined in Topic D.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The first step should be to define "democratize." In the broadest sense, democratization should include providing the NAIRR access to all higher education institutions and associated organizations engaging ML/AI research or applying ML/AI across scientific disciplines. In this context, the limitations to democratize access must include providing the fundamental enabling cyberinfrastructure to those institutions. One envisioning of overcoming these limitations can be found in a recent R&E community-developed paper, "The Minds We Need" (<https://mindsweneed.org/>), which outlines three core actions:

- **Connect every community college, every MSI, and every college and university, including all urban, rural, and tribal institutions** to a world-class and secure R&E infrastructure, with particular attention to institutions that have been chronically underserved;
- **Engage and empower every student and researcher** everywhere with the opportunity to join collaborative environments of the future, recognizing that the "last mind connected may be the mind we need;" and

- **Ensure American competitiveness and leadership** by investing holistically in national R&E infrastructure as a sustainable system.

As highlighted above, particular attention should be paid to institutions that have been chronically underserved. Those institutions have additional challenges, some of which can be found in a recent stakeholder alignment survey conducted by Internet2 and the MS-CC to better understand the needs of HBCUs, TCUs, and HSIs (<https://internet2.edu/solutions/minority-serving-institutions/>).

A key objective of this effort is to help MSIs identify the science, engineering, health, social science, and humanities education and research priorities that call for increased access to and use of data management and computing resources by the higher education community. It is clear that many of the key findings of the survey would apply to the democratization of access to AI R&D.

Key findings of needs identified include:

- **Basic needs:** Across HBCUs, HSIs, and TCUs there is a deep need for basic infrastructure support, such as broadband Wi-Fi on campus and at home for students, staff, and faculty (heightened by the pandemic).
- **Consistency across institutions:** Although there are some unique considerations for certain types of institutions (such as data sovereignty with TCUs), the vast majority of HBCUs, HSIs, and TCUs have similar responses.
- **Workforce development:** Students need literacy and advanced skills with data and computing (the educational mission), while faculty and staff need training and support for a robust cyberinfrastructure (the research mission).
- **Collaboration:** There is strong support for collaboration across institutions to accomplish together what they cannot do separately (with little support for each acting on their own).
- **Institutional operations:** Administrators need a more accessible and responsive data infrastructure for campus operations, surfaced when asked about research data and computing.
- **Societal impact:** There is strong potential for data and computing to advance research on issues central to community culture and disparities in society, with infrastructure as a constraint on achieving these broader impacts.

The first step in overcoming these limitations is to include representatives from underserved institutions and communities from the start of the process.

CONCLUSION

For the foregoing reasons, Internet2 respectfully requests that NSF and OSTP will consider these recommendations and looks forward to working together on this important initiative.

Respectfully submitted,

/s/ Howard Pfeffer
Howard Pfeffer
President and CEO

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Kermit Kubitz

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Subject: RFI Response: National AI Research Resource
Date: Wednesday, September 1, 2021 11:39:51 PM

Artificial intelligence priorities should include:

1. Making sure that people throughout society and academia understand what Artificial Intelligence (AI) and
2. Providing transparency by showing when AI algorithms are being utilized in communications, science, or industry.
3. Comparing results of AI results with human devised systems and results, to see how AI has or has not duplicated or improved the results of human analyses.
4. Identify periodically the results of AI in contributions to various industries.
5. Provide for opportunities for cross-fertilization between and among various AI applications.

Kermit R. Kubitz 

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Lawrence Berkeley Laboratory

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1st, 2021

Via email: NAIRR-responses@nitrd.gov

White House Office of Science and Technology Policy and National Science Foundation

Wendy Wigen, NCO,

2415 Eisenhower Avenue,

Alexandria, VA 22314

RFI Response: National AI Research Resource

Thank you for the opportunity to respond to the National AI Research Resource (NAIRR) RFI. As one of DOE's national laboratories, Berkeley Lab specializes in integrative science and technology, taking advantage of our world-renowned expertise in materials, chemistry, physics, biology, earth and environmental science, mathematics, and computing. We advance the frontiers of science and technology through three approaches: advanced instrumentation and user facilities, large team science, and core research programs led by outstanding investigators.

In 2016 we launched an Artificial Intelligence for Science initiative, which quickly developed into one of our most far-reaching and successful initiatives ever. The initiative took advantage of our strength in applied mathematics and computer science and combined that with our expertise in basic science across all scientific disciplines. The website <http://ml4sci.lbl.gov> describes some of the resulting applications in which AI techniques have been successfully applied to cosmology, particle and nuclear physics, materials science and engineering, chemistry, synthetic biology, genome and biomedical sciences, environmental biology, geoscience, watershed science, climate modeling, technology scale-up, smart grid, water treatment, transportation, advanced detectors, accelerator operations, and many more. Based on our experiences developing and using AI methods and our expertise in managing scientific data through the entire lifecycle, we have the following recommendations.

A National Artificial Intelligence Research Resource must use a holistic approach to combine elements for computation, data lifecycle management, user interface, and training. (Response to Question 1-A)

The stated goal for the NAIRR is to democratize access to the cyberinfrastructure that fuels AI research and development, enabling all of America's diverse AI researchers to fully participate in exploring innovative ideas for advancing AI, including communities, institutions, and regions that have been traditionally underserved—especially with regard to AI research and

related education opportunities. To achieve this goal, the NAIRR must include much more than a set of loosely coordinated computational and data storage resources, but also devote ample attention to providing the tools, technologies, networking, and training to enable researchers to develop and deploy AI methods. Going beyond lip service to the “data lifecycle” to enable truly full-scope lifecycle development from acquisition to preparation to interface is essential. At Berkeley Lab, we have found that multidisciplinary teams composed of experts in AI, user interface design, distributed systems, security, and hardware, in addition to experts in the scientific domain, have been essential to building systems and tools that can push the boundaries of AI research.

- The most impactful AI solutions often require large-scale computing which remains inaccessible for many researchers. New methodologies need to be developed to allow individual researchers to *easily* exploit distributed and heterogeneous compute resources for AI model development. Doing so would democratize large-scale AI, enabling effective use by a broad research community.
- Advanced networking is crucial for democratization of access to key data sets across institutions large and small, urban and rural, and domestic and international.
- There are many barriers preventing researchers from accessing conventional HPC or commercial cloud resources, for example, lack of technical know-how on how to scale from laptop resources to extreme concurrencies, or how to package a workflow into a container, etc., that are not solved simply by providing resources without technical help that will provide on-ramps for researchers with varying levels of skill. The NAIRR needs a plan to provide this.
- There is a need for tools to evaluate the appropriateness and limitations of AI models when applied to research questions, as well as automated model selection and architecture design. Automated model selection will also help address democratization for researchers less well versed in AI. Researchers should also have a credible path to reach out to AI specialists to get their problems addressed.
- It is crucial that the NAIRR allow for models of AI deployment where the workflow is tightly coupled. A key example here is the case of automated experiments. For example, self-driving labs are being proposed for synthetic biology, the operation of beamlines, materials synthesis, earth observations and many others. For synthetic biology, this approach leverages microfluidic chips coupled with real-time DRL models to automate cell DNA modification to improve biofuel yields as part of the “Design Build Test Learn” (DBTL) cycle. The automation can substantially increase exploration of the combinatorially complex state space over current, largely-manual methods, and democratize science by automating a process where human scientific labor is less available.
- The NAIRR should partner with the broader community on leveraging existing and developing additional open standards for AI software and data, and implement these standards for tools developed as part of the NAIRR. This will help ensure researchers can migrate from using the NAIRR to other research-focused resources, or commercial

cloud providers. In addition, standards could include AI benchmarks for evaluating performance on current and emerging architectures.

- As the NAIRR evolves there should be a mechanism for soliciting and incorporating user feedback so that its design and implementation can incorporate diverse perspectives and emerging trends in the field. The NAIRR could also act as a vehicle for user-experience research to improve the overall effectiveness of the resource.

AI software will need to deeply integrate with simulations, data analysis pipelines and legacy codes (Response to Question 1-D and 2-D)

Unlocking the transformative potential of AI in research will require deep coupling of AI and traditional simulation or data analysis applications; as well as in steering and tuning of experimental/simulation workflows with AI. To enable this the NAIRR should recognize:

- R&D will be required to determine both performant and scalable methods for coupling of AI with experiment and simulation, as well as the mix of computational resources to effectively run this mixed workload, including incorporating potential novel AI hardware.
- Research and science communities often have large legacy code bases, written in a mix of programming languages, as well as complex workflows and ad-hoc curation of datasets. Investments should be made to support integration of these with current AI software, such as the python-based deep-learning frameworks.

Data will need to be made *AI-ready* (Response to Question 1-D and 2-D)

The quantity and *quality* of data available to build AI models is a driving determiner of the resulting model's performance. These datasets are dynamic, continually being improved and extended as new data and metadata become available, all the while incorporating expert knowledge. Furthermore, the development of AI models needs to be able to inform the data collection and curation process to enable improvements.

To address these requirements the NAIRR should:

- Federate data centers and partner with the experts that maintain the needed datasets in a shared research infrastructure to enable the AI models to be built using the latest data and expertise.
- Ensure that the needs of AI are incorporated into initial data collection as well as throughout the data lifecycle, including avoiding bias.
- Develop systems that enable tracking of the data and version used in the development of AI systems to allow for retraining and understanding of the limitations of models built from the data.
- Ensure AI model development is part of a cycle of continuous improvement that includes the data collection, processing, and curation needed to drive the models. It will be critical that there is feedback to the data collection and generation process to gather the necessary data and metadata for AI.

- Develop benchmark datasets and challenges that specifically target, and push development of, AI features required for research and science. For example, those that allow for quantifying uncertainty and for the coupling of AI with simulation and data analysis.

Data owners/providers should be able to make data available for AI-based research without requiring full trust of data and computing centers and scientific end-users. (Response to Questions 1-D, 1-F, 2-D, and 2-F)

Two drivers of data owners' reluctance to share data are the risks of sharing sensitive data, and the risks of hosting such data. Even with strong security protections, traditional enclaves still require implicitly full trust in the facility hosting the sensitive data, thereby increasing the liability of an institution for accepting responsibility for hosting data. This limitation can significantly weaken the trust relationships involved in sharing data, particularly when groups are large and distributed. Also, traditional security protections often hinder analysis processes for the scientific community whose abilities and tools are optimized for working in open, collaborative, and distributed environments.

Emerging hardware security technologies, including hardware trusted execution environments (TEEs) can form the basis for platforms that provide strong security benefits while maintaining computational performance, without requiring that system administrators at computing and data centers be fully trusted. Commercial TEEs from the major CPU vendors exist, and have been adopted by the major commercial cloud vendors. To address security concerns of sensitive data, the NAIRR should:

- Leverage TEEs to provide strong security isolation guarantees to protect sensitive data, even from malicious system administrators.
- Support research and development to enable future TEEs to continue to improve both performance and security over today's commercial TEEs to enable a broader range of secure scientific AI applications.

Differential privacy is a statistical technique that can put bounds on the amount of information about a dataset that can be leaked to a data analyst as a result of a query or computation by adding "noise" and enforcing a "privacy budget." It has emerged as an approach to provide strong privacy protection of data output and is now a mainstream solution, with production use by Apple, Google, and the U.S. Census Bureau, the existence of several open source distributions, and successful application to a diverse range of data types. To address privacy and confidentiality concerns of sensitive data, the NAIRR should:

- Leverage differential privacy techniques to enable analysis and AI model training while limiting private information leakage.
- Support research and development to advance the usability of differential privacy and integration of differential privacy in scientific workflows so it can more easily be broadly leveraged for large-scale, scientific AI applications.

Existing Berkeley Lab Resources (Response to Question 4)

- The Department of Energy Office of Science High Performance Computing User Facility NERSC, houses computing and storage resources, including the >120Petaflop Perlmutter supercomputer, with over 6000 NVidia A100 GPUs, a ~120PB Community File System and a 225PB Archival storage system, and data portals for over 8000 Office of Science researchers. NERSC also evaluates next generation AI hardware and testbeds, as well as playing a lead role in developing benchmarks tailored for the research community and for HPC-scale with the MLPerf HPC working group.
- ESnet provides the high-bandwidth, reliable connections that link scientists at national laboratories, universities, and other research institutions, enabling them to collaborate on some of the world's most important scientific challenges including energy, climate science, and the origins of the universe. Funded by the DOE Office of Science, ESnet is managed and operated by the Scientific Networking Division at Lawrence Berkeley National Laboratory. As a nationwide infrastructure and DOE User Facility, ESnet provides scientists with access to unique DOE research facilities and computing resources.
- Berkeley Lab provides expertise and software enabling automated cross DOE SC facility research and collaborations including data analysis from light sources, telescopes, sensors, sequencers and microscopes.
- The lab has demonstrated AI research applications running at increasingly large computing scales, and with increasingly sophisticated approaches (including the 2018 Gordon Bell Prize winner).
- The user facilities and data science platforms at Berkeley Lab - the Joint Genome Institute, the Department of Energy Systems Biology Knowledgebase (KBase), the Advanced Light Source, the Materials Project, and the Molecular Foundry - are premier data generators that fuel new scientific machine learning solutions. In particular, the Materials Project accelerates materials discovery through its comprehensive database of materials properties generated through advanced simulations. The current data set includes >150,000 entries and is used by 150,000 registered users worldwide. The Materials Project has been used to develop many new ML methods, both internally and by the larger research community, for predicting the properties of new chemical compounds.
- NERSC and Berkeley Lab host a number of data sets drawn from experiments in high-energy physics and cosmology, including imaging and spectroscopy for large cosmological surveys seeking to map the effects of dark energy (BOSS, eBOSS, DESI, SN Factory, PTF and ZTF), dark matter search data (LUX, LZ), reactor neutrino data collected over ten years by the Daya Bay experiment, and cosmic microwave background data from more than a dozen ground-based and satellite experiments. We also host the Particle Data Group, a DOE Office of Science Public Reusable Research (PuRe) Data resource, which provides an authoritative, curated database comprising all published

results in particle physics and aspects of cosmology, as well as world averages, uncertainties, and reviews that explain the underlying physics.

- NERSC and Berkeley lab host the ESS-DIVE data repository that contains diverse environmental datasets spanning observational, experimental and modeling research. ESS-DIVE allows data contributors to archive, manage and share various types of data in standardized formats, and obtain digital object identifiers that can be used to cite and track usage of the data. ESS-DIVE users are able to find and obtain data generated by ESS researchers that is organized for better interpretation, analysis, and integration.
- Berkeley Lab runs educational events focussed on the science and research community, including the Deep Learning for Science Summer School and Webinar Series, and practical training for running at HPC scale such as tutorials run by NERSC in collaboration with Intel, Cray/HPE and Nvidia at the Supercomputing conference since 2018.
- Berkeley Lab, with the proximity to University of California and the broader Bay Area has the ability to train students and interns at scale through well-defined internship programs and community outreach. This includes partnerships with the Sustainable Horizons Institute that explicitly seeks to increase representation of students and researchers from marginalized or disadvantaged backgrounds in science.

Public-private partnerships will be the fastest route to resilience and availability for the NAIRR (Response to Question 5)

Over the last 10 years, commercial cloud-computing and storage resources have grown to rival or exceed federally funded high performance computing resources in sheer core count. While HPC resources still often have advantages such as better interconnects and I/O subsystems, the commercial cloud offers resiliency, accessibility, and the ability to surge from using small to larger amounts of resources. As a way to build up capability quickly, a national resource could be composed as an array of compute and data resources, composed of existing high-performance computing resources, or extensions to resources at existing HPC Centers, commercial cloud options, and a layer of software tools. As the resource matures, continued use of commercial cloud resources in addition to dedicated HPC resources makes sense to provide flexibility to cope with surges in demand, and the extensive geographical distribution of cloud resources will additionally provide for more uniform access to data and compute nationally.

In addition, many of the commercial cloud providers have been extremely active in developing artificial intelligence frameworks and other software tools. Often, these frameworks are open-source, and are used by a broad audience of scientists. Here, the NAIRR will need to take care to avoid researchers being locked into proprietary frameworks and datasets that cannot be migrated to other AI resources.

Similarly, multiple datasets have been curated and made available to the wider community by commercial entities with analytical capabilities, e.g. Google Earth Engine, which have been a boon to researchers lacking the resources to do this by themselves. Expanding such resources as part of the NAIRR is another beneficial part of a public-private partnership.

The COVID-19 HPC Consortium could provide a model for how a public-private partnership might function. The Consortium was spun up quickly in response to the demand for computational modeling in the face of the COVID-19 panic, and consisted of cloud and HPC platforms provided by IBM, Microsoft, IBM, NSF HPC Centers, and the DOE National Laboratories. Representatives from each institution formed a governing body to set basic policy, and separate committees were set up to evaluate proposals submitted by researchers applying for access. Successful proposals were directed to the most appropriate resource.

Respectfully submitted,

Deb Agarwal, Division Director, Scientific Data Division

Katie Antypas, Division Deputy, National Energy Research Scientific Computing Division

Wahid Bhimji, Big Data Architect, National Energy Research Scientific Computing Division

Jonathan Carter, Associate Laboratory Director, Computing Sciences Area

Sean Peisert, Staff Scientist, Scientific Data Division

Charuleka Varadharajan, Research Scientist, Earth and Environmental Sciences Area

Peter Zwart, Staff Scientist, Molecular Biophysics and Integrated Bioimaging Division

Contact:

Jonathan Carter



Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Ashik Ghosh

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



August 31, 2021

To

Attn: Wendy Wigen, NCO, 2415 Eisenhower Avenue,
Alexandria, VA 22314, USA.

I am a researcher at the Department of Physics and Astronomy, University of California, Irvine and Physics Division, Lawrence Berkeley National Laboratory, developing AI models that help with our broad spectrum of physics research. I collaborate with the Organisation for Economic Co-operation and Development (OECD) in Paris on topics such as the role of AI in current and future scientific productivity and AI policy, and I am also involved in field of AI Ethics. During my time in France I also informally advised the French computing centre CCIN2P3 (<https://cc.in2p3.fr/en/>) on how to make their AI resources more accessible to researchers in France.

The National Artificial Intelligence Research Resource (NAIRR) is a welcome step to democratize the access to AI research and development infrastructure. Thank you for offering the opportunity to provide comments to the task force. I am responding to the Request for Information (RFI) on an individual capacity. My response to the topics mentioned in the RFI document are given below alongside the topic numbers and letters.

1.

A: In addition to the usual metrics, it is important to measure the diversity of users of the resources. For example the number of research groups (of various academic disciplines) that are embarking on their first AI related project with the help of NAIRR, the fraction of users from smaller research institutions compared to large institutions that already have excellent resources, and the number of individuals without advanced academic degrees who are able to take advantage of the educational resources to start a project. Users from institutions that already have excellent resources would indicate that NAIRR is truly competitive in terms of user-friendliness.

While several small research groups currently maintain their own GPU mini-clusters. This is not efficient because of two reasons. First, it takes significant human time to maintain them for a part-time system administrator. Second, the resources are under-utilized, often free more than 50% of the time. A reliable common national resource may reduce the need to maintain separate mini-clusters. A survey of how many researchers have switched from maintaining their own mini-clusters to using NAIRR would certainly be a useful metric to monitor.

C: While maintaining high utilization for such a large resource requires automated protocols, it must also include the flexibility to request/book for period of guaranteed access to certain resources (such as a single GPU node). This service is offered by the French computing centre CCIN2P3, which is excellent for PhD students at a critical stage of their thesis, before critical conference deadlines and so on.

E: Datasets curated for publicly funded research could be made accessible, in certain cases after an appropriate embargo period.

2.

Educational resources and a software start-up kit targeted at students and researchers with an existing scientific and programming background should be a top priority. There are low-hanging fruits in nearly every scientific domain when it comes to the application of AI. These solutions could improve the state-of-the-art or to tackle a research question in a new way that was computationally infeasible using traditional methods. For many graduate students, the barrier to trying an AI based solution is the large upfront investment of time required to get started. This involves find appropriate learning resources targeted towards researchers

in their particular field and setting up a software pipeline. The time taken for both of these activities are minimized in larger institutions (such as Berkeley Lab) where tutorials and a software start-up kit are readily available. However, for best results, the NAIRR should also host an annual workshop for various researchers to meet in person, exchange ideas, expertise and forge multi-disciplinary collaborations. A small fund (for travel or secondments) could be specifically allocated for collaborations that are forged through the NAIRR workshop. With these measures the NAIRR could boost research productivity and increase the number of innovative ideas that come from smaller research institutions all over the nation.

In addition, educational resources could be made available for people without advanced degrees but with prior programming experience. There are currently far too many application cases for AI compared to the number of individuals with technical know-how. The problems for small communities may not be of immediate interest to large researcher groups or companies but they could be addressed by individuals of that community or organizations working for a social cause, if they have easy access to such resources.

3.

It is absolutely essential for NAIRR to insist that users consciously consider the unintended uses or consequences of their innovations. Warnings may be given along with the publication of material. Educational resources must include components on AI Ethics as well as interpretability and tests for robustness of models. These will be particularly useful for smaller organizations that do not have the resources to employ an in-house AI Ethics expert. A common resource for users who want to understand the ethical implications of their projects may also be provided, either in terms of access to experts or informal discussion forums comprising of other NAIRR users.

6.

AI research evolves much faster than research in traditional academic fields and it is difficult for even the top researchers to keep up with the state-of-the-art. Smaller research groups, particularly those working on applications of AI in their domain of expertise, are unlikely to be able to keep up, despite the support of NAIRR. One possible solution is to host an annual workshop for all NAIRR users with the explicit goal of facilitating the exchange of ideas and forging of inter-institutional, multi-disciplinary collaborations. A small fund (for travel or secondments) could be dedicated to collaborations forged during this workshop.

In closing, I believe the NAIRR is potentially an invaluable initiative to democratize infrastructure that fuels AI innovation. It could help accelerate both small and large research projects and have positive impacts at various levels of society. However, it is essential not only to keep the ethical implications in mind while forming these policies but also to keep them relevant with periodic updates in the future. A successful NAIRR could become a template for other countries and political unions to follow.

Thank you again for providing this opportunity.

Yours Sincerely,

Aishik Ghosh

Postdoctoral Scholar,
University of California, Irvine
CA 92697, USA

&

Affiliate,
Lawrence Berkeley National Laboratory,
1 Cyclotron Rd, Berkeley, CA 94720, USA

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Lawrence Livermore National Laboratory

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

A Development Plan for a National Artificial Intelligence Research Resource

**Lawrence Livermore National Laboratory
Livermore, CA USA**

Technical contacts:

Brian K. Spears

AI Center of Excellence Director

James Brase

Deputy Principal Associate Director for Computing

Chris Clouse

Weapons Simulation Codes Program Director

Robert M. Sharpe

Engineering Deputy Associate Director for Research and Development

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

Executive Summary

The central purposes of the National Artificial Intelligence Research Resource should to focus US AI research strategy and to expand participation in AI research to include the broadest range of researchers possible (roadmap item A). The NAIRR should approach this using four key strategies: (1) the development of a governing consortium drawn from public and private institutions; (2) the identification by this consortium of priority research directions to guide US AI strategy; (3) the production of a set of central products for the research resource comprising common data sets, accessible compute, and common software tools; and (4) the development of partnerships that extend research collaborations between heavily resourced institutions and historically underrepresented ones. (roadmap items C, D, E). The NAIRR should be owned by the Department of Energy owing to its experience with large-scale R&D, interdisciplinary expertise (including AI work and varied applications), and access to critical resources (roadmap item B). The leadership consortium should be led by Lawrence Livermore National Laboratory given its leadership and experience in HPC and AI for applied science (roadmap items B, I)

Need for the research resource (roadmap item A)

The United States needs a coherent, directed AI strategy that both sets priority research directions and democratizes access to the tools and resources needed to advance those research directions. Currently, US AI research advances across economic sectors and national security missions according to a wide set of decoupled goals and motivations. Commercial AI addresses current business and economic needs. Research organizations, especially national laboratories, academic institutions and some commercial and nonprofit organizations, address fundamental methods and national security needs. The efforts of these capable groups need to be further enhanced by concentrating, at least partially, on key national grand challenges – particular AI tasks, methods, and applications found central to a wide range of US competitiveness and security efforts. At the same time, it is critical for future AI research to span

subject matter areas. Key applications – whether in predictive biology, advanced manufacturing, autonomous systems, and more – demand collaboration and interaction among subject matter experts from a wide variety of fields. These include computer science, engineering, physics, biology, chemistry, and applied mathematics. An additional intersection is perhaps the most critical. This is the intersection of simulation with experiment and real-world observation. As an example, at Lawrence Livermore National Laboratory, we have a large strategic effort focused on using AI to merge detailed simulation with high-precision experimental observation. We call this Cognitive Simulation, or CogSim. The ultimate endpoint is to use the combination of simulation and experiment to improve predictive models and accelerate the design of new solutions. However, without acknowledging the coming confluence of data gathered from the edge and experimentation with large-scale simulation, efforts to advance AI methods will miss a large part of the real evolutionary pressure on AI tools. This further highlights the need for a central, organizing body drawn from across the AI R&D ecosystem to identify and set this direction, making sure that efforts draw from all economic sectors, all relevant scientific disciplines, and all sources of useful data (simulation and experiment/observation). The NAIRR should serve as a body to navigate these issues.

The United States also needs a fully engaged and well-trained research and work force to advance the priority research directions. Right now, the lack of access to top tier data sets and computing systems limits full engagement in AI research to large-scale corporations, national Laboratories, and some tier 1 academic research institutions. This excludes engagement by the large talent pool residing in other institutions at tier 1, tier 2, and those focusing on traditionally underserved groups. The US needs the aforementioned NAIRR to build large-scale collaborations across the AI R&D ecosystem focused on developing open data sets, accessible computing resources, and a common public software stack, all of which address the priority research directions established by the NAIRR. The NAIRR needs to enable establishment of critical infrastructure for the dissemination of such data sets to all levels of the AI R&D ecosystem to drive maximal participation. It needs further to enable the establishment of infrastructure for the collation and tracking of subsequent research results, and the disbursement of codified knowledge and best practices across economic sectors and security missions at a variety of skill levels.

In summary, the US needs the NAIRR to build collaboration among the disparate parties in the US AI R&D ecosystem, to identify priority research directions, to develop democratized AI research resources, and to fully engage the totality available US research talent at all levels. The effect of this effort will be to develop a leading AI research capability that outpaces the world by recruiting the effort and abilities of researchers at comprehensive levels not currently available.

Approach to developing the research resource (specific answers to questions 1 and 2 of the RFI)

The NAIRR will dramatically accelerate US AI R&D effectiveness by making available a common research resource. This research resource will be built using a 4-pronged approach:

1. Develop a national consortium built of key constituencies from the US AI R&D ecosystem
2. Identify priority research directions
3. Build (horizontal) partnerships across heavily resourced institutions that produce data sets, compute resources, common software tools to serve as the central products of the research resource.
4. Build (vertical) partnerships and programs that utilize the central products of the research resource to fully recruit R&D talent at all levels for simultaneous advancement of efforts on PRD/grand challenges and development of the national AI workforce and talent reserve.

This approach will develop the central products of the research resource in alignment with national priorities while recruiting a far expanded workforce to drive progress in AI research (figure 1).




Figure 1: a model for directing US AI research, developing shared resources, and fully engaging and expanding research personnel. The NAIRR national consortium provides a cross-disciplinary leadership team guide the evolution and development of the research resource. The consortium identifies priority research directions that capture the dominant US AI needs. Partnerships among the heavily resourced (large companies, national labs, leading academic institutions) provide the central products for the research resource – open data sets, accessible compute resources, and common software tools. Partnerships across all levels use this democratized resource to engage traditionally underserved groups. These partnerships and the research resource combine to open access to AI research while accelerating progress on the priority research directions to advance US competitiveness.

The NAIRR national consortium (partial answer to question 5 of the RFI)

The NAIRR will develop a research consortium composed of key members of the US AI R&D ecosystem. This consortium will include commercial and industrial partners (technology, computer hardware, and computer software companies), national laboratories (national security mission and national resource stewards), and academic partners (drawn from across US academia with emphasis on traditionally underserved groups). The consortium will represent the key constituencies in US AI research. It will draw on this expertise to develop and direct the activities of the NAIRR. The central missions will be: (1) articulating the priority research directions, (2) developing partnerships aimed at producing the shared AI research resources, and (3) developing partnerships that utilize the shared resources, aimed specifically at bridging all levels in the US AI ecosystem with emphasis on historically underserved academic communities, and (4) managing the allocation of research resources.

The consortium should represent the broad range of disciplines and constituencies required for AI research. The disciplines should include expert leaders focused on fundamental AI methods and research as well as those focused on AI method development for applications. The expert leaders should further include those who operate at the intersection of AI development and application, where the applications span science, engineering, and commercial applications.

The expert leaders should be drawn from across the broad range of the US AI R&D ecosystem. This should include parties with substantial resources to contribute, specifically large companies and national laboratories. It must further include academic partners from a wide range of levels – tier 1 research universities, tier 2 schools with fewer resources, and regional schools serving smaller areas and those lacking access to AI resources (data, compute, sophisticated software).

Establishment of this cross-cutting leadership consortium will ensure that a comprehensive set of national interests are represented and that a wide range of resources can be recruited. It will further ensure that participation in the leadership of AI research is spread across all levels of the US AI research ecosystem and that the ecosystem itself is expanded to include those who often lack access.

The priority research directions

The NAIRR consortium will identify key priority research directions or grand challenges. These challenges must span economic sectors and national security missions. They must be curated so that their cross-cutting nature fully envelops the strategic needs of the US in order to maintain international AI competitiveness. They should be posed such that success on these challenges, in the time frames established by the consortium, will clearly lead to US AI global superiority.

The challenges must be interdisciplinary in nature in order to bring the largest field of researchers to the table. They must also drive work that combines computational work, e.g. simulation and modeling, with empirical data (experiment, observation, real-world production). This will drive AI methods that are developed and tailored for interaction and impact in the physical world across economic sectors.

Examples of such challenges might include:

1. Accelerated material discovery: The problem is to discover and optimize molecules to accelerate the development of new materials and medicines to boost economic competitiveness and national security. There has been substantial progress in the application of generative AI models to suggest new concepts with specified molecular properties for small molecules, peptides, and gene sequences, speeding the process of material/bio-material design and discovery. However, these capabilities do not yet exist for complex materials (e.g. polymers) at the same level as for molecular materials, and they do not extend to complex descriptions of materials such as chemical synthesis planning and real-world production. Concerted effort on rapid, AI-driven material discovery pathways will transform US industries that delivery medications, materials, materials synthesis and production methods and more.
2. Enhanced manufacturing competitiveness: The US manufacturing sector is constantly called to deliver a faster, more flexible methodology that will deliver high-quality or novel products in fractions of the current time required. Delivering on this call requires bridging the range of processes in the production chain – material discovery and selection, component and system design, and manufacturing and quality control. We propose a global, AI-driven digital twin framework that spans these processes to shorten delivery times by exploiting recent advances in AI, computing, and manufacturing methods. Research in this area will transform US manufacturing, using AI to produce a competitive advantage in international manufacturing competition that will benefit both commercial and national security operations.

These very brief examples highlight the essential ingredients for meaningful priority research directions. They must engage partners across economic sectors, they must offer wide-ranging benefit to US competitiveness or security, they must develop tools and techniques at the intersection of AI, science, and engineering, and they must operate at an inspirational scale that captures public imagination and future research engagement.

The establishment of these priority research directions by the broad perspectives of the guiding consortium will ensure that a complete set of US AI research needs are addressed by the NAIRR. It will further ensure that the focus and benefits of the resource extend not only across economic sectors but to those groups with previously limited access to AI research opportunities.

The resource production partnerships (partial answer to question 5 of the RFI, roadmap item D)

The NAIRR will further develop multipoint partnerships with the members of the US AI R&D ecosystem who are capable of producing essential resources for the solution of the grand challenges. These essential resources fall into three categories – shareable data sets, accessible compute resources, and common software tools. The production partnerships will feature collaboration between heavily resourced institutions, perhaps a large software company, a large hardware company, and a national laboratory. The collective capabilities of such partnerships will be aimed at making progress on the grand challenges, but equal emphasis will be given to producing **shareable data sets** that others with fewer resources (colleges, regional schools, small businesses) can use to grow their own efforts. The accessibility of such data sets will directly enable all members of the US AI research ecosystem to access the essential data ingredient that is currently unavailable outside of large businesses and national laboratories. The curation of these data sets and the research motivated by them will allow the NAIRR to intentionally focus AI research on areas of national need aligned with the priority research directions.

In similar fashion, production partnerships should be developed to deliver **common software tools**. This will involve developing an AI software ecosystem that spans topics – learning frameworks, workflow software, data tools, and system-level security and management utilities. The software will be interoperable and open source, allowing access by the full spectrum of US AI R&D members, as well as vendor partners, giving them a common target software stack to support. Such a common stack will also provide a uniform interface to AI research, making educational materials universally useful and lowering the barrier to entry for researchers across the AI enterprise. Standardization of interfaces and communication protocols would make it easier to integrate novel AI optimized hardware with more traditional compute resources, widening the aperture for both vendor and researcher participation in the investigation of new hardware approaches. Without common functional tools with a low barrier to initial use and entry, the common data sets and shared compute resources (described next) are functionally useless. Researchers, especially the underserved, will require enormous amounts of time to manage, build, install, and operate software on large computers using large data sets if not equipped with a well-designed and transparent software stack. Note that this will not be a great limiter for the heavily resourced who routinely develop their own tools or purchase others. Thus, a common software tool set will provide an essential key for the underserved to unlock access to other parts of the research resource.

Finally, the NAIRR should develop a **shared compute resource** to democratize access to computational resourced needed to exploit shareable data and common software tools. The NAIRR should draw from the recent experience of the [COVID-19 High-Performance Computing Consortium](#) and on-going efforts to develop a [National Strategic Computing Reserve](#) to support computing in response to national crises. Over 40 U.S. computing organizations, both public and private, came together to provide computing resources free of cost for projects with potential to accelerate understanding and development of treatments and interventions in the pandemic. A similar model could provide a path to providing top-tier computing and data resources to grow the ecosystem of AI R&D in the U.S. This directly addresses question 4 of the RFI. Such a shared computing model provides the final piece that will establish a truly open, accessible, democratized AI research resource. It bears repeating that the resource must include

expanded access to all three fundamental resources – data, software, and compute. Omitting any one of the three will severely limit the set of researchers who can benefit from the NAIRR.

The resource utilization partnerships

The NAIRR will develop a second type of partnership aimed at fully engaging and improving US AI R&D forces at all levels. These partnerships will focus on developing solutions to the directed grand challenges through collaboration across levels. They will be aimed at exploiting the common resources developed by the production partnerships. The efforts should have as endpoints both research solutions and advances, as well as upskilling, re-skilling, and recruitment of researchers. These partnerships must be vertical ones. Some partners should be well-resourced researchers, such as large corporations or national laboratories. They must be paired with historically underserved groups. These pairings are essential democratizing access to AI research. Simply providing underserved groups with the resources (data, software, compute) ignores the fact that progress is made through application of best practices for consuming such resources. To truly reduce the barrier to AI research, underserved institutions need access to these best research practices – methods of teaming, approaches to problems, selection of tools for tasks, optimization of workflows for problem solving, and more. These partnerships will be the final ingredient for expanding the US research ecosystem.

Ownership and Governance of the NAIRR (roadmap items B,C)

The NAIRR should be owned by a governmental agency with experience in large-scale R&D, interdisciplinary expertise (including AI work and varied applications), and access to critical resources (e.g. computing). We recommend that the Department of Energy play such a role. The experience with critical missions, wide-ranging science, world-class computation, and successful partnerships make it a prime candidate for owning the NAIRR. As already mentioned, the NAIRR should be governed by a broad consortium. Such a consortium should have a lead partner capable of steering the group. The lead partner should again have experience with a wide range of computing and science applications, including those pursued through public/private partnerships. The lead partner should also have extensive experience in AI work for a range of applications. We recommend that LLNL serve in this lead partnership capacity. LLNL is strongly coupled to DOE and has a history of global leadership in HPC, having been cited as the most dominant HPC site since 1993 by the TOP500 project. It has a long historical record of successful interdisciplinary science and engineering with a strong track record in public/private partnerships. LLNL pioneered novel risk-sharing relationships that have allowed vendor partners to aggressively innovate; a model which is now widely used across DOE. In addition, LLNL has developed a leading capability in Cognitive Simulation. This CogSim effort combines high-performance computing and simulation with high-precision experimental data using AI methods for applications ranging from drug design to fusion energy to climate modeling to national security. LLNL has created a research and oversight institute, the AI Center of Excellence to help coordinate its wide range of AI research efforts. The AI COE draws from across those efforts to offer a skilled team of technical leads with experience and mandate in developing AI vision and translating that into action via collaboration and R&D. LLNL also has considerable experience in developing and maintaining common software stacks and tools. It is the lead lab in the development of the TOSS operating system, which provides a common operating system for all commodity computing procurements made by NNSA and has made substantial contributions to HPC systems software that are in widespread use across the global HPC community. The history of LLNL, its unique CogSim AI work, and its pre-existing AI COE leadership structure make it uniquely capable of leading the NAIRR consortium.

Existing components (specific answers to question 4 for the RFI)

The DOE and LLNL bring, together with partners, a key set of initial components to help construct the NAIRR according to the plan described above.

Open data sets

1. LLNL provides a large set of open data sets for AI research spanning science missions and data scales. Hosted by the [LLNL Data Science Institute](#) as part of its [Open Data Initiative](#)
2. LLNL has a growing partnership with [MLCommons](#) to further develop open data and benchmarks

Accessible compute resources

1. LLNL has extensive experience from and contributions to [COVID-19 High-Performance Computing Consortium](#). This provides a template for developing the shared compute portion of the NAIRR, as well as a partnership model.
2. LLNL maintains flagship class computing. LLNL machines perennially top the list of most powerful supercomputers in the world as maintained by [TOP500](#).

Public-private partnerships

1. LLNL has an extensive history of successful public-private partnerships and sound partnerships models for doing partnered R&D. [ATOM consortium](#) is a prime example.
2. LLNL has an AI Center of Excellence that expressly develops public-private research partnerships to develop rapid advances in AI research. Their Collaboration Hub model offers a further template for developing NAIRR relationships.

Software tools

1. LLNL has developed and shared a set [Cognitive Simulation](#) tools that lay a foundation for common tools in AI for science ([MaCC](#), [DJINN](#), [LBANN](#), [MERLIN](#))
2. LLNL leads the [RADIUSS](#) project. It aims deploy a common base of foundational scientific software that would be essential to a common NAIRR software stack for use in a shared computing resource.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Lawrence Berkeley National Laboratory Machine Learning Group

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Machine Learning Group
Physics Division
Lawrence Berkeley National Laboratory (LBNL)
1 Cyclotron Road
Berkeley, CA 94702

To whom it may concern,

The goal of the Machine Learning group in the Physics Division at Lawrence Berkeley National Laboratory is to advance the potential for discovery and interdisciplinary collaboration by approaching fundamental physics challenges through the lens of modern machine learning. Our group is a cross-cutting effort that connects researchers developing, adapting, and deploying artificial intelligence (AI) and machine learning (ML) solutions to fundamental physics challenges across the High Energy Physics frontiers, including theory. We are grateful for the opportunity to respond to the Request for Information (RFI) related to the National Artificial Intelligence Research Resource (NAIRR). We have discussed the questions posed to researchers and have some comments and suggestions related to the first question about the roadmap elements that the Task Force should consider.

For the administration of the NAIRR (B.a), we believe it is important to involve both domain scientists and computing/statistics experts. For educational tools and services (D), we think it is critical to support the development of educational resources at the undergraduate and graduate levels. Furthermore, we believe that there should be domain-specific resources. This is particularly important given the large burden of coursework and other training that many students have to take, in addition to the domain-specific requirements. A related topic is the support for common tool development. National Laboratories are particularly well-suited to host software developers for generic and scientific software necessary for AI/ML research and development as well as large-scale deployment (e.g. on High Performance Computing resources). These tools include open source software like scikit-learn, PyTorch, and TensorFlow. In addition to software support, national computing centers could also provide dedicated scientific on demand resources similar to Google's Colab.

The NAIRR specifically mentioned curated datasets (D) and we agree that this is critically important for innovation and development. It is essential that there be physical sciences-oriented datasets, which have unique challenges and thus also unique opportunities for method development that would not be possible with standard datasets like MNIST. In particular, the physical sciences make significant use of large scale *ab initio* simulations, which can become powerful tools for discovery when combined with AI/ML methods. There are also stringent requirements on uncertainty quantification. Additionally, many analysis methods from industry such as anomaly detection do not directly apply. Anomalies in many fundamental physics applications are "group anomalies" where no single example is anomalous and it is only on a statistical basis that one can declare a discovery. Many cutting edge experiments also have stringent requirements on AI/ML inference latency and environment robustness (e.g. to ionizing radiation).

While it is critical to develop and host curated, domain-specific datasets, it is also critical to ensure that there are strategies in place to assess the potential harms of AI/ML applications (G). In particular, there could be hidden costs or biases in AI/ML methods that require domain expertise and/or new AI/ML methodology to mitigate.

Thank you again for the opportunity and we would be happy to expand upon any of the above points.

Sincerely,

Dr. Benjamin Nachman ([REDACTED])
On behalf of the LBNL Physics Division Machine Learning Group
<https://www.physics.lbl.gov/machinelearning>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

**Wayne Gilmore, John Goodhue,
Christopher N. Hill, David Kaelli, Eric
Kolaczyk, Jim Kurose, Scott Yackel**

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1, 2021

Response to “Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource”, RFI 86 FR 39081 (Document number 2021-1566)

As research computing professionals working at member institutions of the Massachusetts Green High Performance Computing Center (MGHPCC), we are pleased to be able to provide comments in response to the above-referenced RFI issued by the National Science Foundation and White House Office of Science & Technology Policy. Our responses here are being made collectively as individuals^[1] with deep experience with research computing infrastructure who collaborate on MGHPCC activities, rather than as institutional responses by our universities or the MGHPCC itself.

Q1: What options should the Task Force consider for any of roadmap elements A through I

Q1.A Goals for establishment and sustainment of a NAIRR

The NAIRR should provide research infrastructure for both *foundational AI* (developing AI theory and methods that are independent of any particular application domain), and *use-inspired AI* research in specific application domains. Such use-inspired research should go beyond simply applying existing AI techniques and add new knowledge and understanding in both foundational AI and use-inspired domains. The virtuous cycle that exists between foundational and use-inspired AI research is a core tenet of NSF’s National AI Research Institutes^[2].

The scope of the NAIRR should be carefully defined, with a clear distinction between funding for development and use of research infrastructure and services (such as the NAIRR) and funding for other aspects of AI research (e.g., faculty, research assistant, staff time; travel). Funding for non-research-infrastructure AI research already exists in numerous Federal agencies and should not be the goal of the NAIRR. The budget for NAIRR should supplement, but not be drawn from, existing agency AI research budgets.

Q1.B A plan for ownership and administration of the National Artificial Intelligence Research Resource

A “non-siloed” infrastructure. The current national computing research infrastructure – funded primarily by the DOE and NSF, but also by NOAA, NIH and other agencies are rather stove-piped systems, with agency-funded resources primarily targeted towards agency-sponsored researchers. As a

national resource targeted at research in a specific area (rather than research funded by a particular agency), the NAIRR should not be partitioned or siloed according to funding agency.

A1.C A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources

Collaborative, multi-agency oversight will be needed for NAIRR. Multiple agencies should have an oversight and governance role within the NAIRR. Agencies with significant investment in AI research (e.g., NSF, DOE, NIH, NIST, NOAA) might be considered as lead agencies in a governance structure. However, since AI techniques are critical in so many application areas, other agencies must also have a seat at the table, in at least an advisory role.

A community-based technical advisory committee. Given the NAIRR will exist to provide critical infrastructure services to the research community, a research-community-based advisory committee should advise the NAIRR on technical and policy issues and direction. In addition, since the results of AI research will impact individuals and organizations far beyond the AI research community, civil society should also participate on this advisory committee.

NAIRR Resource Allocation Committee (NRAC). The NAIRR will certainly be oversubscribed, with resource demands exceeding resource capacity. An allocation process will be needed that grants access based on an assessment of the relative readiness and appropriateness of allocation requests. There are several national models for providing such resource allocation, including NSF's Large Resource Allocations Community (LRAC) and XSEDE Resource Allocations Committee (XRAC), and similar allocation mechanisms for resources funded via the U.S. Department of Energy (e.g., INCITE, NERSC). These merit-based review mechanisms have generally served the community well. One significant shortcoming of the NSF process, however, is that decisions regarding funding for basic research proposals that require HPC resources are made *independently* of decisions regarding the allocation of research computing resources. These two decisions should be coupled, allowing the merit and cost of a research project to be considered as a whole. Certain NASA programs and NSF's CloudBank program can serve as models here, awarding computing resource credits at the time when a research project is reviewed and awarded.

In addition to NAIRR NRAC allocations for specific research projects, NAIRR allocations (and supporting services) will also be needed for educational purposes, for developmental and exploratory projects, and for novel meritorious projects that demonstrate the need and appropriateness for NAIRR resources. In the case of NSF-funded computational resources, a certain percentage of that resource (e.g., 20%) is set-aside for such uses. A community-based committee, perhaps part of the NRAC or perhaps separate, with a makeup that broadly reflects the research community can advise on these allocations and ensure democratized access to NAIRR's AI R&D capabilities.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

NAIRR resources and the public cloud. Whenever possible, NAIRR should leverage commercial cloud offerings, and only develop its own specialized computing resources following a detailed and expansive/inclusive cost/benefit analysis. A 2018 NSF workshop report on *Enabling Computer and Information Science and Engineering Research and Education in the Cloud* ^[3] provides a thoughtful discussion of the possible advantages and disadvantages of this approach. NSF’s Cloudbank project and NIH STRIDES program are important existing programs piloting the notion of cloud-based computation and data resources for the research community (including AI researchers). A valuable question to consider is *can the NAIRR developers envision an NAIRR hosted entirely in commercial clouds - with appropriate high-level support and tailored interfaces for seamless and straightforward use by a wide range of end-users and organizations?*

NAIRR resources should also be able to interoperate seamlessly with campus-based research infrastructure, allowing researchers to migrate code and data among campus-based, NAIRR-based, and commercial-cloud-based services.

The “data resources” hosted by NAIRR should be broadly defined to include higher-level, “synthesized” data products (e.g., knowledge graphs; see “Open Knowledge Network” NITRD Big Data IWG workshop report, November 2018^[4]) and related toolsets. These capabilities are often only available at scale within industry settings; similar open resources and tools will serve to broaden and democratize AI research.

To support the rapidly growing volume of data resources that are envisions to be offered by NAIRR, there will need to be a renewed and increased focus on leveraging new and novel storage technologies and architectures to support this pace. AI workloads are inherently IO-heavy, so will require new architecture and hardware/software solutions to deliver these resources at scale.

Community code (open source, freeware, and purchased) that enhances and manipulates datasets, and open models built on such datasets should be considered part of the NAIRR. The NAIRR should also consider hosting a “model commons” similar to NIH’s notion of a “data commons”.

People. Raw computational capabilities and datasets are a necessary, but far from sufficient, set of resources to meet NAIRR goals. Perhaps most importantly, people-centered services - educational, outreach and training activities, expert consulting, and distributed teams of local, sufficiently resourced, “champions” (perhaps similar to the “campus champions” program that was started by XSEDE and has evolved into community of interest that shares information frequently and widely) - will be key components of an accessible NAIRR that can serve a broad and diverse set of AI researchers.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

A cloud-backed NAIRR provides an opportunity for public private partnership in which the physical infrastructure is provided by commercial cloud vendors. In this case, the true value-added by the

NAIRR is the coordination and funding for resource alignment, access and use, and myriad services layered on top of the physical infrastructure – similar to the approach being taken by Cloudbank, and STRIDES. Multiple cloud providers have collaborated to provide resources to researchers funded under the NSF BIGDATA and Data Hubs programs; commercial cloud services have also hosted NOAA and NASA datasets.

There are a wide array of specific activities that could contribute to a NAIRR network and that do exist today. The NAIRR planners may want to consider strategically collaborating with, strengthening and leveraging some/all of these entities, rather than creating a duplicate structure. As an illustrative example we list below a non-exhaustive set of projects/activities that a NAIRR effort could help bring into a more coordinated whole.

1. All commercial cloud providers have public dataset programs that are curated and organized collections with practical value for applied AI. The Azure "Planetary Computer" initiative is one (of many) examples of the sort of remarkable resources available this way.
2. Commercial cloud providers also support low-powered, free to use compute and software resources (for example Google colab) that can be powerful platforms for basic educational infrastructure.
3. Commercial cloud providers also support access to large scale compute resources and to ad-hoc data storage for a fee.

All these commercial resources are a tremendous potential resource to leverage in some way. There are some caveats to these resources that a NAIRR could try and improve.

- Public datasets (that are freely hosted) are selected by cloud providers and held under terms that are set by providers. The eligibility and storage duration of a potential public dataset is not guaranteed.
- Datasets are subject to metered egress charges between providers and/or between providers and external networks (including academic networks). No provider has a mechanism to unconditionally waive all metered egress charges.
- Cloud-based compute costs can be sub-optimal for some modes of use, an issue that presents challenges for both academic and industry users^[5].
- Efficient use of commercial cloud resources requires cost engineering and financial management skills that (today) rarely the academic department-level, where researchers are accustomed to submitting allocation proposals for resources that are operated by agencies such as NSF, DOE, NIH, or by campus research computing groups.

- The commercial providers offer powerful but proprietary AI tool sets. On one hand these tool sets accelerate scientific progress; but on the other hand, they serve as obstacles to the goal of being able to move flexibly from one provider to another. The NAIRR could play a role in ensuring that these proprietary differences represent healthy competition of ideas, and not artificial barriers (e.g., via proprietary data formats).
- General cloud use is treated as services and subject to F&A under current Federal sponsor practices. The Financial Accounting Standards Board^[6] has recently made some moves toward capitalization allowances for some cloud expenses. The NAIRR planners should carefully consider the myriad issues around how F&A is charged for cloud services.

4. In the non-commercial space there are many exemplar infrastructure activities that an NAIRR could build on. Some of these include

- NIH STRIDES and Bridge2AI- these programs have been successful in improving data availability across the broader NIH community
- NSF Open Storage Network - this pilot program has goals to provide a managed storage fabric for very large data collections that have value but may not be eligible for cloud provider public dataset hosting (for example very large corpuses or collections for which the science audience is valuable but small in number).
- The Dataverse community - a mature network of curated long-term data storage for research reproducibility and extension with a strong heritage in social and economic research data.
- NSF CloudBank - a streamlined program for allocating cloud resources within grant allocations

5. In the human resources space some example infrastructure activities NAIRR could leverage and strengthen include the:

- NSF Campus Champions, Cyberteam and ask.ci network - these activities support training and knowledge sharing across technical domains
- NSF XSEDE extended support program - this provides technology expertise to help domain projects leverage advanced computing well
- Northeastern University led AI Jumpstart initiative aimed at proactively training small and medium scale industries in AI leveraging state-of-the-art technology solutions.
- MIT and Harvard Airforce AI Accelerator programs aimed at connecting cutting edge AI research with strategic national security needs
- The DARPA Colosseum 5G simulator at Northeastern, a platform to enable AI researchers to explore applications of AI to next-generation wireless.
- BU SAIL program aimed at embedding AI researchers within domain research groups
- US Research Software Engineer Association network. This is a self-organizing group of research software engineering professionals, including technical AI/ML practitioners in academia, that is working to build a professional national community focussed on sustaining careers of people who sustain software infrastructure that underlies all manner of research, including AI and ML.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

We believe that deliberate public-private partnering has excellent potential for creating an infrastructure that meets the broad national needs that a NAIRR initiative ought to aspire to, if it is to play a major role in national research and economic growth. We have discussed a number of public-private partnerships above.

The COVID19 HPC Consortium may provide aspects of an interesting exemplar on the potential of a collaborative multi-stakeholder public-private grouping that works effectively as a whole. Although the Consortium came about in very unusual circumstances, the activity demonstrated an entity that spanned NASA, DOE, NIH, NSF and DoD, universities, commercial cloud participants AWS, Azure, Google Compute, and IBM. The participants developed effective mechanisms for researchers from all sectors (industry, academic and federal) in need of resources to access a diverse mix of capabilities that included compute, data, and technical expertise all within a single overarching virtual organization. While the exceptional circumstances meant that important issues around financial and business bookkeeping were largely set aside, many other practical governance and operational issues that a close public-private NAIRR partnership might need to manage were addressed effectively. The NAIRR planners might want to leverage the lessons learned from this activity on how to reach a broad research community, provide relatively seamless access to resources, and proactively support and guide projects to the correct scale and type of resources.

A significant public-private partnering around NAIRR would be very much in keeping with more deliberate Federal, industry, and academia coordinated innovation strategies envisioned in the PCAST report “Recommendations for Strengthening American Leadership in Industries of the Future”^[7].

^[1] Contributors to this document are *Wayne Gilmore*, Director of Research Computing Services, Information Services & Technology, Boston U.;; *John Goodhue*, Executive Director, Massachusetts Green High Performance Computing Center; *Christopher N. Hill*, Principal Research Scientist, MIT; *David Kaeli*, College of Engineering Distinguished Professor of Electrical and Computer Engineering, Northeastern U.; *Eric Kolaczyk*, Professor of Mathematics and Statistics, Director of the Hariri Institute for Computing and Computational Science & Engineering, Boston U.; *Jim Kurose*, Distinguished Professor of Computer Science and Associate Chancellor for Partnerships and Innovation, U. Massachusetts Amherst; *Scott Yockel*, Director for Research Computing, Faculty of Arts and Science, Harvard U.

^[2] <https://www.nsf.gov/pubs/2020/nsf20604/nsf20604.htm>

^[3] <https://cra.org/cloud-access-for-nsf-cise-research/>

^[4] <https://www.nitrd.gov/pubs/open-knowledge-network-workshop-report-2018.pdf>

^[5] <https://a16z.com/2021/05/27/cost-of-cloud-paradox-market-cap-cloud-lifecycle-scale-growth-repatriation-optimization/>

^[6] <https://www.fasb.org/>

^[7] https://science.osti.gov/-/media/_/pdf/about/pcast/202006/PCAST_June_2020_Report.pdf

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Mathematica

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Akira Bell
SVP & Chief Information Officer

October 1, 2021

Wendy Wigen

Re: RFI Response: National AI Research Resource

Dear Wendy Wigen:

Thank you for giving Mathematica and other members of the Artificial Intelligence (AI) research community an opportunity to provide input on issues being considered by the National Artificial Intelligence Research Resource (NAIRR) Task Force. A Research Resource is a logical next step to build on the Evidence Act recommendations and infrastructure created by the Open Data initiative. Current cloud technology capabilities can provide the ability to store, analyze, and visualize massive amounts of data. Although AI methods are becoming more accessible, they are not yet ubiquitous and this rapidly changing field has not yet established methods to address equity issues.

For more than 50 years, Mathematica has been at the forefront of assessing the effectiveness of policies and programs to improve public well-being. Our deep bench of more than 1,500 experts translates big questions into insights for our public and private sector partners. We apply our expertise at the intersection of data science and social science by leveraging data assets using advanced technologies and methods such as AI, reusable and scalable data and platforms, and high-performance secure cloud infrastructure. We have a reputation for quality and objectivity rooted in our rigor and commitment to improving well-being for all, which centers on the ethical use of data and equity issues. We offer recommendations for the NAIRR based on our expertise as end users of data and methodologists who have a deep understanding of government policies and programs.

In our experience, good AI requires strong interconnections between the data, technology, and methods. Understanding of the data content, program issues, and questions decision makers need answered are often missing links that cause AI projects to fail. Mathematica's approach to AI provides actionable recommendations by grounding our work in an understanding of the policy issues, pairing the expertise of multidisciplinary subject matter experts with methodological experts, and using data ethically with an emphasis on equity at every stage.

As such, I am pleased to represent my colleagues in our response to **Questions 1A, 1Bii, 1D-1G, 2, 3, 4, and 6**, as published in *Federal Register* no. 2021-15660 by the National Science Foundation and Office of Science and Technology Policy. We share these insights in the hope they help to inform your work and facilitate important discussions and action among federal agencies.

Please direct all follow-up questions or comments to my colleague David Roberts, Mathematica's director of Strategic Communications.

Q1: What options should the Task Force consider and why?

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

The NAIRR seeks to increase access to computational resources, data, educational tools, and user support for AI research. As such, goals should include the successful adoption, use, and sustainability of this resource, which requires effective promotion, establishing appropriate use guidelines, and implementing metrics to track usage. To incorporate key indicators as appropriate, the Task Force should define metrics for measuring success as the resource is designed.

The Task Force should develop a strategically targeted communication and outreach strategy to raise awareness among potential users, particularly those disadvantaged historically with respect to resource access to address equity. The NAIRR should require site registration and include tracking metrics for specific resources (data sets, training, and so on) to collect data that demonstrate where uptake has been most successful. The NAIRR can in turn use this information to target communication messages to achieve the objective of supporting a broad range of AI researchers.

To use resources effectively, the NAIRR should promote an approach to AI research founded in methodological rigor, including input from topical subject matter experts (SMEs) and designed with end users in mind. SMEs serve as a critical bridge between the data, technology, methods, and good AI systems by ensuring AI researchers and system developers understand the issues the AI system aims to solve. Similarly, incorporating a human-centered design (HCD) approach will help ensure that end users can practically implement resulting AI products. The NAIRR should also provide guidelines for practitioners related to the responsible use of AI products, such as ways to evaluate equity and bias in the use of AI products.

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including: ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

Governance of the NAIRR should include government and private sector organizations to leverage knowledge of government processes and structures with expertise in cutting-edge applications of AI to social science problems. This mix of participants, along with time limits for governance committee members, will help ensure committees remain current on AI industry trends and represent a range of perspectives. Selection criteria for private sector organizations should include companies engaged in AI with public and private sector lines of work across various industries. The NAIRR should also establish safeguards, such as term limits or conflict of interest agreements, to ensure private organizations do not participate for their own gain.

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

To maximize access and use among AI researchers, the NAIRR should be a cloud-based portal with a range of computational and analytic resources curated by topic and tailored by experience level—from beginner to more highly trained AI practitioner. Mathematica found this approach was successful for increasing access to our [COVID-19 Curated Data, Modeling, and Policy Resources](#).

Organizing the breadth of resources and providing ongoing support will require extensive curation and coordination. The NAIRR should establish robust and scalable metadata standards to help organize and maintain resources, such as [Document, Discover and Interoperate](#) standards. Using open-source tools can help control cost, and they have the benefit of substantial community support.

At a minimum, the NAIRR should include several core capabilities:

- Ingesting, cataloging, storing, and archiving data sets
- Version control for data sets and other resources
- Computational resources and visualization tools for working with data
- Configurable environments with the ability to preserve them for future use or save documentation about the configuration
- Monitoring and analysis of system access, resource use, and system security
- User access controls
- Metrics to track system access and use of specific resources by individual users
- Communication tools researchers and SMEs can use to discuss specific data sets or resources either privately or publicly (with discussions accessible to the NAIRR community)

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

Researcher knowledge, access to appropriate tools and resources, and agency publication of high quality data sets can all be significant barriers to disseminating and using NAIRR data sets. To understand and address these issues, the NAIRR should provide a variety of resources:

Access to SMEs. SMEs provide critical context for data sets and how to practically address problems using AI. Providing access to SMEs who are knowledgeable about data sets hosted in the NAIRR and related policy issues, as well as supports for communicating with SMEs, will help ensure researchers understand and properly interpret data sets used in AI research.

Support for multiple data formats. Providing data sets in multiple formats or conversion tools can increase accessibility among researchers who have limited experience with data manipulation tools or programming languages. AI researchers can also have a broad interest in using government records that are publicly available, but stored as unstructured data sets (for example, emails, *Congressional Record*, and other text-based data). To facilitate the use of unstructured data sets, the NAIRR should provide support for NoSQL databases or data lakes and resources to provision unstructured data sets or structured databases in third normal form (that is, tidy format) whenever possible. The NAIRR could also include semi-automated text analytic tools to facilitate exploring these resources or reference existing supports for the use of unstructured data.¹

Tools for managing data access. Some NAIRR data sets might be static, but others could reside elsewhere or be available only through live data streams. Providing support to access live data streams or data sets housed outside the NAIRR, particularly using application programming

¹ Existing examples of resources to support the use of unstructured data include <https://www.ibm.com/cloud/blog/structured-vs-unstructured-data> and <https://www.essentialsql.com/database-normalization/>.

interfaces (APIs), is critical and can provide the portal an advantage over other methods of accessing data. The NAIRR can apply best practices for data sharing as outlined by the [HHS CDO initiative](#).

Data matching guidance. Different data sets might code similar constructs inconsistently and provide hard-to-follow data documentation, making it difficult for users to combine data sources. Promoting standards in data use guides for NAIRR data sets and providing guidelines for combining data sets can help address usage barriers related to users' knowledge. Encouraging standard, detailed data set documentation can also alert AI researchers to potential issues related to responsible use and equity associated with data quality issues often introduced during processing.

A data operations framework. A framework (such as DataOps) can define an approach for designing and implementing the NAIRR portal in addition to outlining best practices, recommended workflows, and use of guidelines for AI researchers. Providing a framework such as this can support access to data for use in AI research by removing the requirement for researchers to have the knowledge to develop and implement workflows and follow best practices on their own.

Lessons learned from data-sharing initiatives. The NAIRR should consider lessons learned from existing data-sharing and open data initiatives about what works well and potential pitfalls. For example, incorporating incentives for sharing data, providing support for agencies to share data in requested ways, developing channels for data users to provide feedback (such as ratings and data use popularities) and suggestions on shared data, and establishing clear lines of responsibility and ownership are all important.

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;

The Task Force must consider requirements such as security of the NAIRR portal, integrity of the resources housed there, and user-level access controls. The Task Force should tailor management strategies based on the level of risk to balance security with accessibility. These strategies should align with [federal thinking on cybersecurity](#) while providing flexibility and ease of access for the broad range of AI researchers the resource is intended to support. A [Zero Trust Architecture](#) approach best suits this type of resource: the focus must be on the data and the level of protection required based on the risk profile of the data, and the purposes the data serve. Real-time, contextual details that explicitly describe each user should be the basis for granting access. For example, the potential risk associated with access to public data is low and should not require substantial resources. By contrast, the integrity of the data sets and guidance documentation, as well as availability of computational resources, pose a higher level of risk that warrants greater investment.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

Privacy, civil rights, and civil liberties inherently intertwine in AI research. Multiple stages of the research can introduce privacy and bias issues—from selecting to cleaning data and developing and implementing the model. Public use datasets are reviewed for risk of individual disclosure and often remove personal identifiers such as race, gender, and ethnicity. However, combining data sets from disparate sources could result in reidentification of individuals based on new field combinations. The NAIRR can holistically assess the data sets it hosts to ensure possible end users cannot maliciously or inadvertently reidentify individuals by combining data sets. [Standards for deidentifying data sets](#) provided for AI research might need to be more stringent. The NAIRR should also establish a disclosure risk committee to evaluate potential risk based on the combination of data sources

available. It might also consider a record identification system that can facilitate matching across sources and allow for removing fields that could result in reidentification.

To address bias that infringes on civil rights or civil liberties, the NAIRR should provide guidance to researchers on how to address algorithmic bias, such as a process to evaluate how resources might reinforce inequity and provide responsible use guidance to researchers. This guidance should focus on transparency in algorithms so researchers can understand what the models do and how they reach conclusions. If data sets available on the NAIRR include sensitive data such as race or ethnicity that could potentially lead to infringement of civil liberties, the NAIRR should implement more stringent requirements for adhering to provided guidelines.

Q2: Which capabilities and services provided through the NAIRR should be prioritized?

The goal of establishing a national computing and collaboration environment for AI researchers and students is ambitious and timely. To realize this vision, the NAIRR should prioritize several factors:

Scalable, simple, and low-cost infrastructure. There are many options for relatively cheap, secure, easily provisioned cloud-based infrastructure with free tiers for students and researchers to use. The NAIRR could provide access to Amazon Web Services (AWS), Azure, and/or Google Cloud Platform resources directly through the portal, with subsidies or expanded free tiers for students and researchers whose research serves a public mission. This would enable the NAIRR to be stood up more quickly and allow AI researchers to take full advantage of cloud-based services.

Access to large data sets. The NAIRR should provide general access to curated public use files (PUF), as well as access to restricted use files (RUF) with secure controls. It should also support standardized APIs that authorized end users can query programmatically so they can seamlessly integrate their services and data sets into existing data pipelines. Although the NAIRR should incentivize open access to data, proper vetting of third-party API providers is essential to reduce the risk of malicious or unauthorized use or provision of data. Potential users should present compelling evidence for why they need access to APIs with sensitive data. In addition, the NAIRR could provide access to larger real-world and so-called toy data sets researchers could use to follow along with training materials to help students and aspiring AI researchers learn big data processing and analysis techniques.

Coordination with SMEs. The NAIRR should provide access to knowledgeable SMEs who can help students and researchers understand and properly interpret data sets. These SMEs should include agency personnel with deep expertise in NAIRR data sets, as well as outside volunteers with domain expertise. Experienced AI researchers can provide insight on the analytical approaches, and SMEs can help aspiring and seasoned data scientists identify domain-specific problems, understand the intricacies of complex or idiosyncratic data sets, and ensure models are developed with end users in mind. The NAIRR should also provide communication channels—through video conference software, Slack channels, communities of practice, or other virtual communication mechanisms—to enable users of the platform to reach out to SMEs and confirm they properly understand and interpret their findings. Unfortunately, curating data sets and making SMEs available can be costly, time consuming, and resource intensive. To address these resources constraints, the NAIRR should give certain users of the platform—such as students and researchers working on high-priority government contracts—privileged access to SMEs and expedited access to curated data sets.

Technical documentation, training, and tutorials. The NAIRR should pair access to data sets and SMEs with technical documentation, tutorials, and standardized metadata to empower users. The portal should include (English and non-English) materials and tutorials written for multiple

free-to-use programming languages—including Python, R, Julia, C#, and Java. In addition, the portal’s design should integrate with popular version control and containerization services (such as GitHub, Bitbucket, Docker, and Kubernetes). Adoption of the NAIRR is likely to be higher the more it prioritizes popular open-source programming languages and tools. Using legacy languages or not integrating with tools commonly used in the industry will lower credibility and push people toward other platforms.

Human-centered design. Researchers often view AI projects as a technology or data science effort aimed at achieving a certain level of accuracy or technical rigor, rather than an effort to improve human decision making. However, to be valuable, practitioners must adopt and use AI to drive improvement and change. As part of its mission, the NAIRR should promote incorporating HCD at the earliest stages of an AI project to properly frame analytic problems; develop solutions that deliver results to the right people at the right time; and provide wraparound support in the form of explainability, transparency, and user support. Mathematica’s work on the [Centers for Medicare & Medicaid Services Artificial Intelligence Health Outcomes Challenge](#) reflects this mindset. Our team implemented an HCD process with multiple rounds of discovery, research, and user testing with patients, doctors, and other clinicians to accomplish these goals. The result was analytic output that could be explained and interpreted easily by clinical end users, as well as actionability by integrating with established workflows to create a seamless user experience that turned outputs into solutions.

Q3: How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

The NAIRR and its components should establish clear and actionable ethics standards, document potential bias and nuance in resource manuals, set guidelines to enhance appropriate use of the resource, and protect data privacy. Specifically, the NAIRR should consider the following aspects to reinforce ethical and responsible AI research and development:

Create measurable and reportable ethics standards and rubrics. When developing ethics standards and rubrics, the NAIRR can increase awareness and discussion of work done by advocacy groups (such as the [Algorithmic Justice League](#)), media outlets (for example, [ProPublica](#)), academic institutions, and other organizations focusing on responsible use of AI and algorithmic bias. Equipping interactive AI ethics tools in the ethics rubrics will make the standards more actionable. For example, Mathematica developed an [open-source tool](#) that enables users to assess fairness in algorithms predicting binary outcomes and quantify the uncertainty of each estimate under a Bayesian framework.

Assess potential bias in NAIRR data sets and models and provide considerations in use documentation. The NAIRR should assess the data sets it makes available through the portal for underlying biases, and it should clearly articulate any potential biases to potential end users. Sometimes it might be difficult to discern the types of biases that underly a given data set, especially when working with large or unstructured data. For instance, natural language models trained on text data might reflect the unconscious biases of the original authors. Furthermore, the NAIRR should assess predictive models developed using NAIRR for their proxy discrimination before applying to other applications within the portal. Complex predictive models can often obfuscate prejudiced conclusions by using features that might not seem to correlate with race, ethnicity, and so on, but actually do. AI developers should assess their models for proxy discrimination when validating their results. The NAIRR should require documentation of such nuances in the data user manual and

keep updating the document to incorporate feedback from data users on underlying data biases through a crowdsourcing mechanism (such as a user-reported rating and feedback system).

Require proposals to address and account for potential bias when developing AI tools using NAIRR resources. The proposal should require researchers to assess biases that might affect the data sets they want to use to develop a new AI tool and outline their plan to account for these biases when developing the new tool. For example, if a developer planned to use Medicare claims data to develop an AI prediction tool, it would be important to consider regional variations in health care service use and the underlying racial biases in health care delivery, both of which could introduce substantial bias into prediction tools. AI developers should also consider whether implementing or using the tool they propose to develop could adversely affect historically disadvantaged priority groups. For tools with the potential to adversely affect certain groups, it will be important to assess the actual impact on those subgroups when they implement the tool. When validating a new AI tool, developers should provide information on the tool's performance for relevant subgroups. For example, developers should provide AI model performance metrics, such as the Area Under Curve (AUC) score for subgroups based on gender, race and ethnicity, and age, among other individual characteristics.

Establish data report agreement to enhance transparency in AI development. If AI developers receive access to a data set or multiple data sets for the purpose of developing an AI tool, the developers should have to provide information regarding how representative their training, testing, and validation data sets are of the population to which they hope to apply the AI tool. They should present such information for data sets directly used for AI model training, testing, and validation, rather than the original data sets from the NAIRR. For example, if a developer excludes certain variables because of data quality issues or other reasons in the process of cleaning the data to build an AI tool, the final cleaned data might not be representative even if the initial data were.

Deidentify data when possible to ensure disclosure avoidance. The governance and oversight structure for the NAIRR must account for different types of data that impose different levels of risk to people or groups represented by those data. Because it is unlikely obtaining consent for use of particularly sensitive data—such as health or financial data—will be feasible, the NAIRR should deidentify these data when possible. Moreover, the NAIRR should carefully consider data privacy issues that are introduced when providing access to many data sets. For example, when using multiple linked data sets, researchers must still protect the confidentiality of deidentified information, checking by differentially private algorithms to ensure disclosure avoidance.

Provide disclosure guidelines or requirements for AI tools developed using NAIRR resources. AI developers should make enough information available about an AI tool to enable a potential user to determine whether it is appropriate to use in a specific population or setting, to understand the performance of the tool, and to understand how to interpret prediction results. Such information includes the data used to develop the AI tool, the approach to develop and validate the tool, how the algorithm reaches its results, and information regarding overall AI model performance and its performance for specific subpopulations. Transparency does not always require that AI developers make an AI algorithm itself available though. The level of transparency should depend to some extent on the level of risk associated with the use of the AI tool, the level of precision of the predictions, the clarity of the recommended actions to end users, and the potential for legal liability ([NAM AI report](#), pp. 192, 219–220).

Q4: What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

The NAIRR can benefit from government open data initiatives, evolving interoperable and digital information, technology offerings from government contractors and cloud service providers, and established private–academic partnerships. The NAIRR can use several types of building blocks:

Existing open data initiatives from the federal government and agencies. For example, [Data.gov](#) includes more than 300,000 data sets and rich data tools, data incubator sources, and skills development resources that enable citizen participation in government, create opportunities for economic development, and inform decision making in both private and government sectors. The [Data Optimization Initiative](#), led by the Office of the Chief Technology Officer within the U.S. Department of Health and Human Services (HHS), has released more than 4,500 data sets collected by HHS for public use by researchers and entrepreneurs.

Health-related AI research that incorporates evolving interoperable and digital information. HHS is in the process of advancing the connectivity of electronic health information and interoperability of health information. For example, the electronic health information exchange enables doctors, nurses, pharmacists, other health care providers and patients to appropriately access and securely share a patient’s vital medical information electronically, improving the speed, quality, safety, and cost of caring for patients. Future health-related AI research developed using the NAIRR could apply and incorporate information from such data resources.

Existing technology offerings and data sources from government contractors. Many government contractors and research institutes have invested in technology offerings and curated data sources that could support specific needs of NAIRR users. Specifically, many of them operate virtual data enclaves (VDEs) that provide secure access to restricted data through a virtual private network connection to a portal on researchers’ computers. For example, the data library recently developed by Mathematica is an AWS cloud solution that enables users to search for, discover, and interact with ingested data sets in a secure and compliant environment. NORC offers a data enclave that provides data storage and management, computational resources, and reporting tools. The NAIRR could also benefit from existing, curated data sources. For example, Mathematica surveyed and organized a [repository of publicly available data, modeling, and policy resources on COVID-19](#) that enables states, health care decision makers, providers, and others to predict need and direct resources based on the best available evidence during the pandemic.

Strong private-academic partnerships. Partnerships with academic institutions that foster AI research can expand the reach of the NAIRR and build a pipeline of future contributors. For example, the [Howard-Mathematica SICSS partnership](#) is an instructional program hosted at Howard University, a historically black college and university (HBCU), designed to promote learning and support development of expertise in computational science for graduate students, postdoctoral researchers, and beginning faculty from HBCUs and underrepresented communities. This program highlights how innovative partnerships can help counter anti-Black racism and inequity. The program includes lectures, group problem sets, participant-led research projects, office hour sessions led by industry professionals, and invited outside speakers who conduct computational social science research in a variety of settings. Topics covered included text as data, website scraping, digital field experiments, machine learning, and ethics. Partnerships such as this introduce new views on existing applications of AI that those who are not members of underrepresented communities might not always clearly understand, such as limitations of natural language processing for African America

vernacular English, AI tools for health care practitioners to eliminate bias for rural or unrepresented populations, and AI studies of network effects for immigrants.

Cloud service provider collaboration. NAIRR can work with cloud computing providers for flexible and scalable resources, while benefitting from the volume of services offered through these providers. Many organizations already partner with these cloud providers to develop AI solutions to address social needs and improve public well-being. For example, [Mathematica is collaborating with Google Cloud](#) on Google Cloud’s new Healthcare Data Engine to provide health care measures, data analytics, and data science capabilities to health care teams and leaders across the country.

Q6: Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The past decade has seen a dramatic increase in private sector investment in AI research and development, both in the United States and throughout the world.² These commercial investments have often had a democratizing effect on AI research, particularly through corporate participation in open-source and open data communities. However, many challenges remain to ensure AI resources are widely available and researchers with varying experience levels have the capabilities to develop and deploy AI. Among the core obstacles to broader democratization of AI resources are access to (1) adequate training and expertise; (2) large data sets, especially private or industry-specific data sets;³ and (3) high-performance computing infrastructure necessary to use many state-of-the-art AI algorithms and data processing techniques.

Access to training and expertise. AI research often requires advanced knowledge of computer science, as well as a strong foundation in mathematics such as calculus, probability and statistics, and linear algebra. This can create barriers to entry for aspiring AI researchers or those coming from fields other than science, technology, engineering, or math (STEM) who are less well versed in these concepts. If not paired with adequate training, educational resources, and subject matter expertise, AI democratization can lead to inaccurate statistical interpretations, developing biased models, and programming or technical errors that produce incorrect results. The NAIRR should provide training resources that emphasize developing data science and computer science skill sets, proper data preparation and modeling procedures, and accurate interpretation of model performance and results. A central goal of the NAIRR’s training efforts should be to reach a broader and more diverse talent pool of prospective AI researchers coming from colleges and universities, technical institutes, and other non-traditional learning settings. To advance this goal, the NAIRR could also consider establishing certificate programs for aspiring technologists that would be more easily attainable for those with fewer resources at their disposal. Finally, the NAIRR should provide channels of communication with SMEs who can support students and researchers on an as-needed basis.

Access to data sets. Beyond training and expertise, access to data—particularly large volumes of data needed to train highly accurate machine learning models—is essential. Although private sector investment has helped to advance many AI tools and procedures, there are often limits to sharing data in the private sector. A key to democratizing AI research for the NAIRR will be providing easy access to large data sets while maintaining proper data security and role-based authentication

² Zachary Arnold. “Tracking AI Investment Initial Findings from the Private Markets.” September 2020. Available at <https://cset.georgetown.edu/publication/tracking-ai-investment/>.

³ Ahmed, Nur, and Muntasir Wahed. “The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research.” *ArXiv*, October 2020. Available at <https://arxiv.org/abs/2010.15581>.

procedures. Ideally, the NAIRR platform would provide access to both government and private sector data and provide tools for greater integration between the two.

Access to computing infrastructure. Another core challenge to democratizing AI research is unequal access to the computing resources required to process large volumes of data. NAIRR can address this issue by providing tools and infrastructure tailored to experience levels, along with clear guidance and user documentation. Tools should be available at little or no cost to ensure accessibility, but the NAIRR must develop a prioritization and approval process to ensure specific users do not hoard resources or use them for unlawful purposes. In addition, the NAIRR will have to provide a help desk support function to respond to users' issues and requests.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Medical Imaging and Resource Center, University of Chicago

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

86 FR 39081 Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource

Submitted by Medical Imaging and Resource Center (MIDRC)

Q: Roadmap for a shared research infrastructure that would provide Artificial Intelligence (AI) researchers and students across scientific disciplines with access to computational resources, high-quality data, educational tools, and user support.

A: Many deep-learning AI architectures are available as open-source (on GitHub, for example), often including weights from pre-training on the ImageNet dataset. Thus, it is relatively simple for anyone to start developing an AI tool. It isn't as easy, however, to perform a study that has potential clinical utility. The major hurdles to democratizing AI research are the availability of 1) reliable, ethically-sourced, trustworthy data, 2) computational resources, 3) education in proper study design, and 4) performance evaluation tools and services (including sequestered data).

We, as the Medical Imaging and Resource Center (MIDRC; midrc.org), are responding to the RFI by presenting our roadmap, which we believe can be extended to benefit the scientific medical disciplines. We are a multi-institutional initiative currently developing new diagnostics, including AI/ML algorithms, and comprised of researchers from academia, imaging societies, private practices, government agencies, and industry. As such, we are able to leverage each collaborators' existing infrastructure, processes and expertise within the medical imaging community, all of which already meet accepted standards of quality, security, access and sustainability. We serve as a linked data commons that harmonizes access to data and data management activities across all participating organizations at three critical stages: 1) intake, including curation, de-identification, harmonization, and quality assessment, 2) annotation and labelling of images and other data using semi-automated approaches, and 3) distributed access and query methods.

In our domain of radiology/medical image analysis, high-quality public medical image repositories (data commons) are an extremely valuable resource. Several exist in the realm of cancer imaging, such as the Cancer Imaging Archive and the Imaging Data Commons, with the former also including datasets with genetic information available in the Cancer Genome Atlas. At MIDRC, our developing data commons is aimed at becoming the world's largest highest-quality repository for COVID-19 imaging and metadata, with the framework built to extend to imaging of other diseases. Reducing the barriers to access to high-quality, de-identified, curated, annotated data levels the playing field, not only for training algorithms, but also for benchmark testing of different algorithms' performance to enable translation through regulatory to benefit the public good. To date, MIDRC has released over 13,000 imaging studies to the public, with another 50,000 currently undergoing quality control prior to release; with currently 23 clinical sites in the contributing to releasing pipeline. MIDRC is establishing functionality for researchers to search the data commons and create diverse cohorts, with which trustworthy AI/ML algorithms can be developed. MIDRC is also creating a sequestered

commons to serve as an independent source of data for assessing the performance level of developed AI/ML algorithms, expediting the regulatory approval process. MIDRC is also conducting joint pilot studies in collaboration with non-imaging COVID-19 data repositories, to link imaging and non-imaging data of overlapping patient cohorts, allowing the full wealth of data available to be more efficiently and fully explored.

Researchers without access to the computational resources required for AI applications (such as GPUs), or without the capability to store large amounts of data locally, would benefit immensely if imaging data repositories were coupled with a computational enclave for development and evaluation of AI tools. Such a computational enclave is currently in the planning for MIDRC as we begin collaborations with a national lab and various cloud computing companies to synergistically utilize their extensive compute power and services.

Another focus of MIDRC is research and user-education. We are actively developing user resources to guide in proper study design and evaluation. MIDRC has quarterly Town Hall meetings and a monthly Seminar Series, allowing members of the medical community at-large to interact with our investigators as they discuss their current research and noteworthy advances. We release regular updates on our work through newsletters and press releases, and provide full support to both data contributors and end users of the data through the Contact Us feature on our fully-operational website.

Publicly available tools are also in development to educate users of our data on methods to properly design their studies and evaluate the performance of their algorithms, such as an interactive decision tree-type interface that guides users to various analyses methods, software tools, literature, and performance metrics, based on their task (such as detection, diagnosis, segmentation, time-to-event analysis, estimation), reference standard, and type of AI output.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

S. Joseph Sirintrapun

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



Memorial Sloan Kettering
Cancer Center

8/31/2021

Attention:

Wendy Wigen, NCO
2415 Eisenhower Avenue
Alexandria, VA 22314

Re: Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource (NAIRR)

To whom this may concern:

As Director of Pathology Informatics and the current President of the Association for Pathology Informatics (API), I continually engage with digital and AI technologies along with their integration in healthcare. I share my personal opinions in response to the RFI on an implementation plan for the National AI Research Resource (NAIRR) in this letter.

Listed in the subsequent pages, I provide my responses for the Task Force to consider to all roadmap question items (A through I), underlined, and my reasoning. As part of my responses, I also give some opinions on how NAIRR and its components can reinforce principles of ethical and responsible research and the development of AI. Finally, I also include building blocks for the NAIRR, namely government, academic, professional organizations, and industry.

Again, I appreciate your attention.

Sincerely,

S. Joseph Sirintrapun, M.D.

Director of Pathology Informatics
Associate Attending
Department of Pathology
1275 York Avenue Rm A515, New York, NY 10065

Memorial Sloan Kettering Cancer Center

www.mskcc.org

NCI-designated Comprehensive Cancer Center

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

My overarching goals are for 1) democratizing availability of healthcare data to AI community stakeholders, 2) sustaining a collaborative community environment that brings together all the engaged stakeholders for continual AI development, and 3) ensuring that clinical AI deployment is effective and ensures trust in its use

Gaps:

- No systematic study is underway for a sufficient level of “anonymization” to share data between institutions to aggregate for AI development
- No systematic research is underway to understand how to capture, organize automatically, and “index” healthcare data at scale, with much of the progress in capturing data through inefficient manual curation
- No practical/usable ontologic knowledge frameworks that engage both human domain experts and machines, and likewise, there are no systematic means to develop AI tools that work off these frameworks
- No strategic direction on the regulation of AI and how to handle “change control” or versioning of algorithms as they evolve
- No incentive to “operationalize” AI into clinical workflows with process engineering that will ensure trust in matching the AI intended to use to the clinical situation at the correct point in time.

Addressing gaps through funding and resource allocation directed towards:

- Creation of innovative and sustainable incentive models (i.e., ONC, HHS/CMS) for health institutions to anonymize, curate, and release healthcare data to cloud infrastructures
- Study the science of anonymization of different healthcare data types and how to make such data portable
- Study the art and science of anonymization for different kinds of healthcare data, and then build standardized tools that scale and are extensible
- Study the art and science of evaluating demographic and social determinant data, and then produce data models by which tools get built to capture, qualify, and quantitate biases and disparities
- Study the art and science of multimodal diagnostic data (i.e., clinical impression,

radiologic, pathology, and laboratory), and then build data models that enable the feature and label engineering for AI tools

- Study the art and science of patient status and clinical outcome multimodal data (i.e., clinical impression, radiologic, pathology, and laboratory), and then build data models that enable the feature and label engineering for AI tools
- Study “change control,” or the handling in how updated versions of AI get deployed into clinical environments.
- Study of implementation science with AI, meaning how to deliver AI into clinical workflows; this means:
 - Deploying AI at the right time and situation with minimal friction
 - Optimized trust that the AI intended use case meets the clinical situation at hand in the correct context when triggered
 - Maintaining the right balance of “human in the loop” not to overburden yet ensure governance over AI tools to not becomes too autonomous without trusted oversight

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and

The solution is not just one agency but a cross-agency partnership between HHS entities, including the FDA, HHS, and ONC. This partnership includes professional clinical organizations with AI expertise (i.e., The Association for Pathology Informatics, a member of the FDA Network of Experts). Also included in the partnerships are industry and big tech stakeholders, all in a collaborative environment to discuss the issues and solutions to overcome gaps and even provide resources.

ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

A governance structure will require a cross-agency partnership between HHS entities, including FDA, HHS, and ONC, and in collaboration with professional clinical organizations with AI expertise (i.e., The Association for Pathology Informatics, which is also a member of the FDA Network of Experts). Also included in the partnerships are industry and big tech groups, in a collaborative environment to discuss the issues and solutions to overcome gaps and even provide resources.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

No such model exists and will need iteration as the cross-agency partnership with stakeholder engagement evolves.

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

May require engagement by the big tech cloud platforms (i.e. Amazon, Google) about computing infrastructures.

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

Bringing the multiple stakeholders together through coordination, leadership, and commitment with funding with issues that arise through the duration of the effort; all are requirements to ensure the sustainability of this effort.

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;

No comment, since I'm not a security expert. The science and "operationalization" of anonymization is critical in sharing data without exposure to sensitive information.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

No comment, since I'm not a privacy/ethics expert. The science and "operationalization" of anonymization is critical in sharing data without exposure to sensitive information.

H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; and

May need big tech to participate through a "win-win" incentive partnership model for such a significant government/academic/vendor collaboration. This effort will require leadership with a solid strategic vision.

I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

No comment, only to state that such a resource should ensure that academic/professional organizations should get incentivized for sustained engagement in the collaborations.

1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

See my responses above.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

See my responses for A above, where I list the gaps and my opinions on funding and resource allocation.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

- Study the science of anonymization of different healthcare data types and how to make such data portable
- Study the art and science of anonymization for different kinds of healthcare data, and then build standardized tools that scale and are extensible
- Study the art and science of evaluating demographic and social determinant data, and then produce data models by which tools get built to capture, qualify, and quantitate biases and disparities
- Study and develop better models in data use that ensures trust and ethical use for the generators of the data

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

One building block is leveraging partnerships with existing professional organizations like the Association for Pathology Informatics (API). API is the foremost organization with internationally recognized top experts in the field of pathology informatics. In addition, API has joined the FDA Network of Experts and the Network of Digital Health Experts. The FDA Network of Experts consists of a select group of 100 professional non-profit status organizations that identify relevant member experts and connect them to the FDA program members with timely feedback (not opinion or recommendation) on newly emergent health technologies. API, having experts in digital pathology and AI and synchronizing with the FDA Network of Experts mission, is a valuable existing resource for NAIRR.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Such partnerships are integral in advancing and incentivizing future sustained thriving development. One exemplar is the Defense Advanced Research Projects Agency (DARPA), seeking potential public-private partnerships to further the Robotic Servicing of Geosynchronous Satellites program.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

See my responses for A above, where I list the gaps and my opinions on funding and resource allocation.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Microsoft

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

**Response of Microsoft Corporation to
White House Office of Science and Technology Policy
and National Science Foundation RFI on
the National Artificial Intelligence Research Resource**

1st October 2021

Introduction

Microsoft enthusiastically supports the aims of the National AI Research Resource (NAIRR) Task Force and shares the vision of democratizing access to the evolving ecosystem of foundational computing capabilities to empower a larger and more diverse artificial intelligence (AI) research and development (R&D) community. In addition to the foundational role that compute cycles play in this R&D ecosystem, critical assets such as datasets, advanced AI tools, simulation environments, and platform models are playing an increasing role as key enablers of AI R&D.

As a global technology company and member of the larger R&D ecosystem, we have three primary recommendations for consideration by the NAIRR Task Force:

Recommendation 1: Leverage the rapidly evolving landscape of computing capabilities, advances, and practices available through commercial cloud platforms

Over the last several decades, we have experienced an increasing interdependence between advances in AI R&D and the development of AI platforms and services, with each enabled by rapidly evolving large-scale computing infrastructure, AI system software, and large-scale datasets as depicted in Figure 1. At Microsoft, AI R&D has been increasingly enabled by the co-development and use of our own AI research platforms and supported by intensive computation provided by our global network of data and compute centers.

We see the empowerment of a collaborative, global research community with expanded and democratized access to secure cloud-based compute, scalable AI infrastructure, and advanced AI platforms as both consistent with NAIRR goals D, G and H and aligned with Microsoft's mission. Microsoft research scientists and engineers have broadly and consistently engaged with the global research community through both direct collaboration and public-private partnerships. Microsoft's research and engineering teams maintain a commitment to open science and have co-authored and shared thousands of publications, many following from rich collaborations with researchers around the world. In addition to direct collaborations, Microsoft has also engaged in public-private partnerships in multiple countries including U.S.-based programs such as the COVID-19 HPC Consortium¹, The National Science Foundation's Cloudbank² program, and the National Institute of Health's STRIDES³ program. These programs demonstrate the potential for effective partnerships across the public and private sectors to support the advancement of AI R&D.

Recommendation 2: Create a resource framework that offers large-scale infrastructure, AI system software, data, and platform models to support the AI R&D community

Consistent with NAIRR goals D, F and H, a NAIRR solution should enable a broad range of research workflows spanning many different disciplines and areas of study and support effective, efficient, and

secure execution of the associated computing workloads. Just as software development has come a long way from assembly language, computing resources that support scientific research computing are evolving into higher-level building blocks and tools. While large-scale and high-performance infrastructure provides the lowest common denominator for advancing research computing, higher-level tools in the AI technology stack (e.g., as depicted in Figure 1) aim to automate essential elements of current and anticipated future workflows relevant to scientific research. Given the heterogenous landscape of AI services, this broad range of workflows will be best supported by leveraging services from multiple providers that have the capacity to offer these capabilities for research. Such a resource would benefit the AI R&D community by enabling broad access to ongoing innovation, unique breakthroughs, and distinctive services being created across the public and private sectors. The need to increase access to these services has been illustrated by the recent emergence of large pre-trained neural models as platforms for both AI research and application development. A research paradigm has become central in AI R&D where fixed large-scale *platform models* are adapted via “fine-tuning” procedures to develop capabilities for specific tasks.

Commercial cloud capabilities, augmenting on-premises infrastructure, could enable a robust resource framework with minimal or no queue times and allocation limits for users. Historically, major technology shifts have occurred every 2-4 years as evidenced by recent advancements in networking, security and encryption technologies, and personal assistants. Modern commercial cloud capabilities offer the potential to enable significant advancements across the AI technology stack to occur on the order of months. A NAIRR solution should have sufficient agility to support the incorporation of new AI advancements as they become available for use. The ability for the broader research community to keep pace with these advancements will, in turn, help accelerate innovation across the AI technology stack.

Recommendation 3: Harness the creation of a National AI Research Resource to advance U.S. workforce development goals

The U.S. government has articulated a need to significantly ramp up the capacity to develop and skill a diverse next generation of AI researchers and engineers⁴. In reference to the NAIRR goals D and H, we believe that industry can play an important role in advancing educational tools and services, and overall, AI workforce readiness. Shared infrastructure between industrial and academic research, such as in Microsoft’s AI and IoT Insider Labs⁵, will improve collaboration across sectors in advancing technology to address societal challenges at scale while enhancing the transparency and reproducibility of research breakthroughs and strengthening the research to production pipeline. Academia can leverage the industry connection to enhance relevant curricula⁶ with state-of-the-art computing capabilities that are already ubiquitous in industry. By engaging industry, including large corporations, small and medium technology companies, and early-stage startups, the responsibility of catalyzing robust research to innovation pipeline will be broadly shared.

The remainder of the document provides further detailed ideas on how we can accomplish these three recommendations in collaboration with the broader research community.




Figure 1. AI Technology Stack, illustrated by representative Microsoft services^{13,14,20,24,25}

Recommendation 1: Leverage the rapidly evolving landscape of computing capabilities, advances, and practices available through commercial cloud platforms

With reference to Question 2: *Which capabilities and services provided through the NAIRR should be prioritized?*

It was recently stated by panelists at the Computing Research Association (CRA) led Virtual Roundtable on Best Practices on using the Cloud for Computing Research⁷ that research created in education institutions will not keep up with 21st century advances if it doesn't take advantage of the enormous capacity, rich software, and hardware infrastructure that the commercial cloud offers. The NAIRR should accommodate a variety of research and simulation workloads (as illustrated in Figure 2) across systems, platforms, and resources. In general, research computing spans a very diverse set of workloads including core topics in AI and broader computer science as well as in a broad spectrum of areas across sciences

and engineering. Topics of particular interest cross a broad swath of frontier technologies and research with distinct compute needs. As examples, finite element analyses for computational fluid dynamics typically require high memory and core while climate modeling typically requires ensemble calculations. Given the ubiquity of AI across many areas of research, a robust NAIRR solution is likely to require similar breadth and diversity.

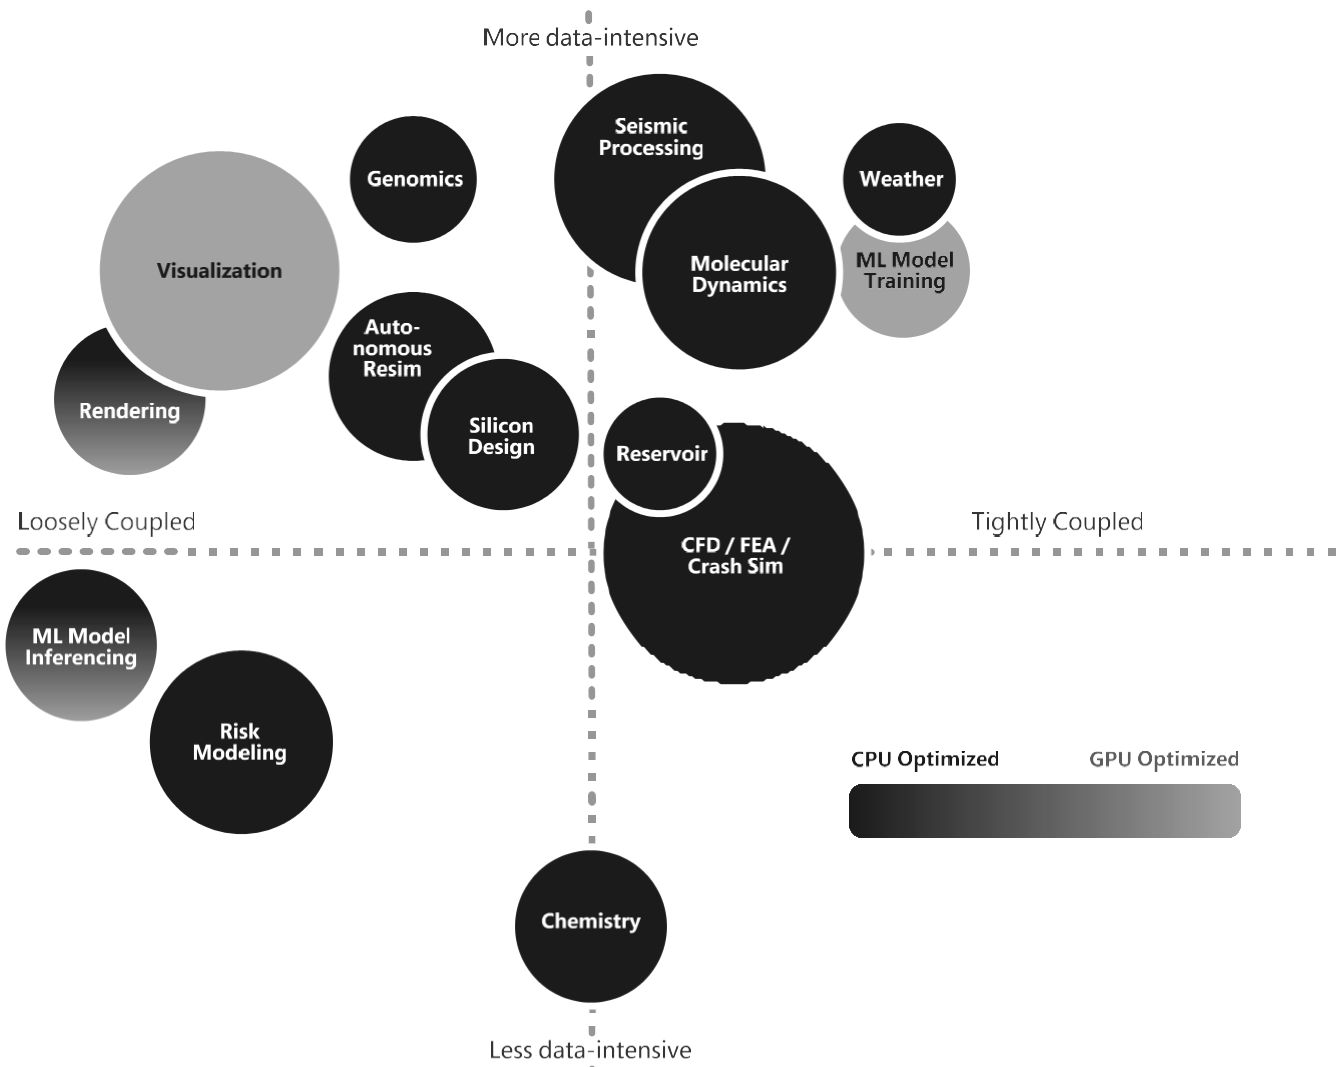


Figure 2. Cloud Compute Workload Mapping²⁶

With reference to Question 4: *What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?*

Depending upon the specific computing need, a single subscription to a commercial cloud account could enable a researcher to access a broad range of compute resources optimized for performance and efficiency for the relevant workloads. For example, Azure’s VM instances with InfiniBand-enabled clusters provide significant performance advantages for running tightly coupled high-performance compute workloads. Figure 3 provides an ‘at a glance’ reference to the Azure VM families. In addition, VMs that incorporate

confidential compute capabilities offer the ability to build secure enclave-based applications within trusted execution environments.

As climate change and health concerns continue to impact society, they represent an increasing influence on computing needs associated with government-funded research. Hence, the workloads and datasets that support these research areas are expanding beyond traditional modeling and simulation.

Azure can be connected to datasets hosted in federal agencies and if needed, secured by a Virtual Private Network tunnel. Azure ExpressRoute is a service that enables users to create private connections between the datacenters that cloud servers are hosted in, and infrastructure that are on premises or in a co-location environment with dedicated circuit available from 50 Mbps to a 100 Gbps pair.

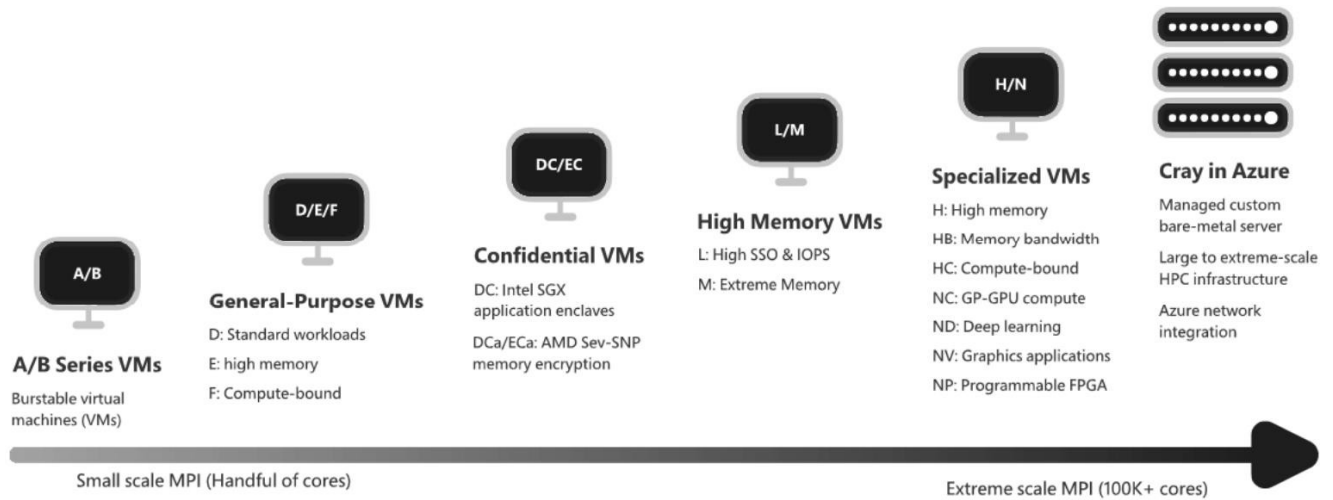


Figure 3. Azure Virtual Machine Groups²⁷

Given the increasing scale of commercial cloud, cybersecurity represents a growing need and area of continuous innovation and advancement. For instance, Azure ingests trillions of security signals each day and combines security, compliance, identity, and management as an interdependent whole. The application of AI to the analysis of these signals enhances the ability to identify, detect, protect, and respond in real-time to known and emerging cybersecurity threats.

Azure is accredited²² at FedRAMP Moderate and High, DoD Impact Levels (IL) 2,4,5, and 6, and meets ITAR, CJIS, DFARS, and NIST 800-171 requirements. Additionally, Microsoft complies with NIST SP 800-53 and is aligned with NIST SP 800-161 supplemental guidance.

With reference to Question 5: *What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?*

NAIRR/NSF-funded researchers should ideally have sufficient understanding of requirements associated with relevant workflows so that they are equipped to evaluate or take full advantage of a cloud-supported NAIRR solution. NAIRR should work toward a model that facilitates this level of understanding to help ensure that each workflow runs with high efficiency on readily available resources.

In a cloud-computing setting, each scientific workflow is, in effect, its own benchmark. Benchmarking the workflow in the target environment enables in-situ tuning and optimization with respect to

performance and cost. While individualized benchmarks take time on the front-end, they eliminate the benchmark-to-execution step and can inform a framework of pre-computed simulation environments.

Programs such as NSF Cloudbank and NIH Strides represent emerging efforts leveraging public-private partnerships to promote a culture of government-funded research supported by cloud compute capabilities. Outside the US, recent examples include the UK government’s billion-dollar investment through which the UK Met Office and Microsoft partnered to build the world’s most powerful weather and climate forecasting supercomputer,¹⁶ the establishment of a government-funded public-private 66,000m² AI innovation hub in China¹⁷, and the projects undertaken by the Swedish National Center for AI in partnership with the private sector¹⁸.

Recommendation 2: Create a resource framework that offers large-scale infrastructure, AI system software, data, and platform models to support the AI R&D community

The Fourth Industrial Revolution has resulted in widespread application of machine learning and AI technologies across a broad range of industrial automation use-cases (eg. as depicted in Figure 4). Emerging capabilities like Github Copilot⁹ are accelerating the development and application of new capabilities by enabling developers to automatically identify relevant code written by other developers with the same intent, code search systems can help automatically retrieve relevant code based on natural language queries. This AI-enabled code intelligence has been the intent behind open benchmark datasets like CodeXGLUE¹⁰ and developer tool enhancements powered by Github coPilot that aim to empower the 23 million+ developer community.

AI and machine learning aren't new concepts, and many of the theories have been unchanged for decades, but recent technological advances have accelerated AI innovation. These advances include large-scale




Figure 4. AI Use-Cases in Industry⁸

infrastructure such as AI supercomputing¹¹, and advanced storage and bandwidth capabilities to enable more accurate and useful algorithm predictions. There is a clear need for higher-level AI services beyond the CS and industrial automation community to computational scientists and other advanced practitioners.

In reference to Question 2: *Which capabilities and services provided through the NAIRR should be prioritized?*

Recognizing that the NAIRR intends to serve communities across a diversity of scientific disciplines, it is important to acknowledge the need for higher-level AI services to facilitate access to AI resources by research communities. As illustrated in Figure 1, these resources fall under the following broad categories:

A. AI System Software

Training a large neural model requires several components such as deep learning optimization strategies, efficient and scalable runtime engine, and a service to manage the experiments and the distributed training infrastructure. Workflows include model design, distributed training, mixed precision, gradient accumulation, and checkpointing.

ONNX Runtime is an open-source engine to support the highly efficient high performant training and inference of AI models in a framework-agnostic manner across a range of hardware. This runtime brings together efficient implementations of the mathematical operations underpinning the deep learning algorithms and the training optimizations from various Microsoft capabilities into one integrated package.

Azure Machine Learning (Azure ML) is the end-to-end machine learning development lifecycle that enables efficient building, training, and deployment of machine-learned models at scale. Azure ML enables team collaboration with experiment tracking, model performance metrics collection, industry leading MLOps, i.e., DevOps for machine learning. Azure ML Service supports model training and deployment across all major deep learning frameworks and runtimes including the ONNX runtime, leveraging the Azure AI infrastructure including large GPU clusters.

Today we're seeing the most sweeping changes in data management since the relational database revolution of the '70s and '80s. These advances motivate significant changes in how researchers engage with the next generation of data management platforms.

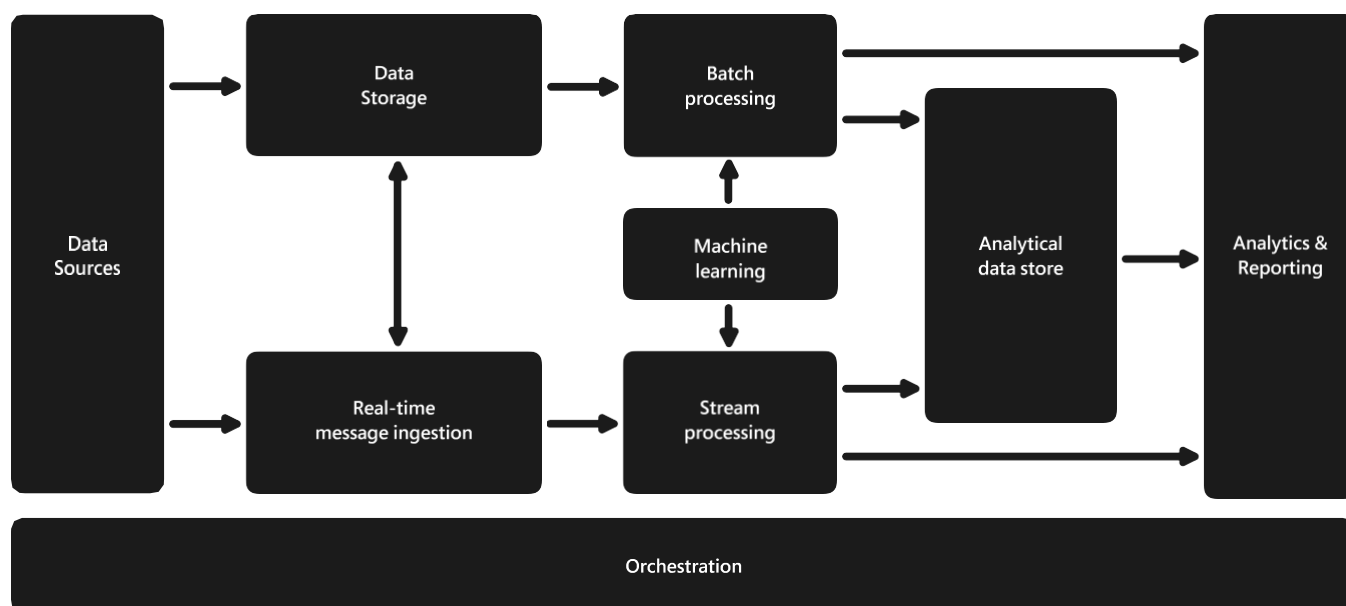


Figure 5. Cloud based Modern Data Platform Reference Architecture²³

The components of Microsoft's Modern Data Platform such as shown in Figure 5 are designed to address such demands.

A modern data platform architecture is designed to handle ingestion, storage, processing, and analysis of data that is too large or complex for traditional database systems. In addition to size (volume), the types of data (variety), speed of the data (velocity) and inconsistencies in the data (veracity) is also considered. Azure supports several "Big Data" deployment architectures, to include the Lambda (hot and cold data processing) and Kappa (single path stream processing) architectures which provide a consistent way to ask questions of data in motion or data at rest. Azure also supports multiple other products build to simplify the management of AI/ML processes such as Data Science Virtual Machine (DSVM), Azure HDInsight, Azure Databricks, Azure Synapse Analytics and Azure Machine Learning Service (Azure ML).

B. Models that have been pre-trained with massive data

Following the trend that larger natural language models lead to better results, we are building upon the strong collaboration across research and engineering to achieve new modelling techniques for language breakthroughs. Language generation models are currently used in many natural language scenarios such as word prediction, sentence completion, text summarization and are an important research area in the future of AI. Microsoft developed a family of large-scale natural language processing models, a few of which are listed in Figure 1, and in 2020, announced Turing model¹³ for natural language generation (NLG), the largest model ever published at 17 billion parameters, outperforming the state of the art on a variety of language modeling benchmarks while accomplishing numerous practical tasks such as summarization and question answering.

The deep learning library behind the model, DeepSpeed was made open source to make distributed training of large models easier. DeepSpeed contains the Zero Redundancy Optimizer (ZeRO) for training models with 100 million parameters or more at scale. Both DeepSpeed and ZeRO are available to researchers, because training large networks like those that utilize the Transformer architecture can be expensive and can encounter issues at scale.

Through the Turing Academic program¹² Microsoft provided access to the Turing family of models, to support general and open research, including efforts aimed at advancing principles of learning and reasoning, exploring novel applications, and pursuing better understanding of challenges and opportunities regarding the ethical and responsible use of large-scale neural language models.

The initial results of the program have resulted in research collaborations with several academic research groups that explore research areas across a broad range of AI research problems such as to 1) transfer learning in medical notes 2) understand and quantify biases in models 3) identify linguistic markers that exacerbate bias 4) distinguish between real and fake events.

The program has shown promise as something that can have far reaching impact in examining the unintended impacts of language model-based AI technologies. Eventually the language model approach of how AI is developed from narrow, custom models to multi-purpose, massive models is expected to be generalized to multi-modal text, video, voice data. Research investment via a convening of public and private sector stakeholders is critical to balance the economic incentives that make the resultant AI technologies inevitable in use by commercial interests.

In the exploration of language models, as well as other forms of AI, it is important to consider the responsible design, development, and deployment of those technologies. Overall, there is growing investment²⁰ on the responsible development and fielding of AI systems, including developing

accountability and governance across industry. Efforts on responsible AI innovation are focused on how to operationalize the following principles:

- **Fairness.** AI systems should treat all people fairly
- **Reliability & Safety.** AI systems should perform reliably and safely
- **Privacy & Security.** AI systems should be secure and respect privacy
- **Inclusiveness.** AI systems should empower everyone and engage people
- **Transparency.** AI systems should be understandable
- **Accountability.** People should be accountable for AI systems

Moving forward it will be important to seek scalable, research-based methods to advance these principles, such as outlined in the NSCAI report¹⁴ that presents key considerations for the responsible development of AI. Measurement and evaluation tools to assess the technical requirements for responsible AI, for example areas like fairness and accuracy, will contribute to this goal.

In addition, several AI benchmark datasets and models are such as MS-MARCO for large scale reading comprehension and question answering, MT-DNN for natural language understanding, XGLUE for cross-lingual evaluation benchmark XGLUE are available under the AI for Scale research¹⁵ effort.

Recommendation 3: Utilize the creation of a national AI research resource to advance U.S. workforce development goals

With reference to Question 5: *What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?* And question 6. *Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?*

The research communities' experiences over the past two decades indicates that AI research is currently dependent on *both* pioneering researchers and advanced engineers. We expect this situation will continue through the next few decades of research in AI. In our AI research efforts and in working with academic researchers in AI, we have periodically seen a limitation to the pace of research can be the researchers' ability to employ state-of-art computing systems and techniques. These are techniques that are enabled by collaborating engineers, and in contrast, the researchers who are able to rapidly and efficiently leverage their computing resources have some advantages over peers who do not.

Academia provides a unique and foundational educational environment that prepares individuals for a lifetime of learning. It develops the emerging workforce by fostering their abilities to think critically, to grow techniques and capacity to learn; to approach new problems creatively, and to understand and evaluate conflicting ideas and information. In doing so, academia creates the knowledge foundations required for most researchers and engineers that are responsible for the next wave of AI innovation in the US. Academic learning content and materials take a relatively long period of time to develop, evolve in an iteratively manner, and are vetted by peers and communities of educators.

We also recognize that undergraduate and graduate curriculums have more content and material than they can practically accommodate in their degree programs. And we realize that it is not appropriate to re-orient degree programs towards courses and work dedicated to the development of skills on specific tools and platforms. However, to reach their full potential throughout their career, researchers and engineers do need to develop some explicit skills to understand and employ the continually changing AI systems, techniques, and tools. These are skills that Industry is well-equipped to provide.

It will also take a combined effort of academia, industry, and the government to increase the diversity of AI researchers and engineers. We can only accomplish this through sustained initiatives that impact students through their educational journey and professionals through their careers.

NAIRR should embrace the critical, complementary roles that both academia and industry play in educating and skilling researchers and engineers. Building, operating, and using the next evolution of compute and data resources will require foundationally sound and continually skilled workforce. Together, academia and industry will provide institutional capacity and commitment to adequately staff, in both quality and quantity, the workforce essential to research workflows.

Microsoft has a strong history for supporting the development of future-ready skills in collaboration with the academic community through co-curricular programs like Microsoft's AI Business School and Microsoft Learn for Educators. AI Business School²¹ (AIBS) is a master class curriculum designed to build knowledge and confidence in AI. It provides comprehensive training for non-technical learners spanning critical topics like strategy, culture, responsible AI, scale AI, AI for business users, and AI technology for leaders. Its content has been adopted widely by colleges and universities across the globe.

Our experiences with Microsoft Learn for Educators¹⁹ (MSLE) is anchored in the belief that higher education intuitions and faculty members play a pivotal role in empowering students for future success. As such, MSLE supports students, faculty members, and higher education institutions with free curriculum, training, and tools for teaching technical skills. MSLE supports faculty at colleges, universities, and community colleges who want to help build their students' skills in technical topics like cloud computing, AI, data engineering and security, so that they can attain industry-recognized certifications that prepare them for their future careers. Microsoft Learn for Educators currently trains and provides technical curriculum to over 3500 faculty members at 450 colleges, universities, community colleges, and polytechnics across 85 countries and in 9 languages. In fiscal year 2021 alone this Microsoft program helped higher education institutions provide cutting-edge technical skilling learning experiences to over 60,000 students.

Summary

Advancing AI in a manner that is trustworthy, ethical and benefits the whole of society will require participation and collaboration across a wide range of scientific disciplines, institutions, and sectors. Microsoft is aligned with the NAIRR task force vision of democratizing access to the cyberinfrastructure that fuels AI research and development and looks forward to continuing to participate in upcoming forums related to the mission.

References

¹ <https://www.hpcwire.com/2021/03/26/the-covid-19-hpc-consortium-looks-ahead-to-a-national-strategic-computing-reserve/>

² <https://www.cloudbank.org/about>

³ <https://www.nih.gov/news-events/news-releases/nih-expands-biomedical-research-cloud-microsoft-azure>

⁴ <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>

⁵ <https://www.microsoftinsiderlabs.com/>

⁶ <https://www.businesswire.com/news/home/20210114005727/en/>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

NSF AI Institute for Artificial Intelligence and Fundamental Interactions

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

NSF IAIFI Response to National AI Research Resource Task Force

The NSF AI Institute for Artificial Intelligence and Fundamental Interactions (IAIFI, <http://iaifi.org/>) is supportive of the national AI Research Resource (NAIRR) task force effort as available computing resources are vital to the success of our Institute and of other AI research endeavors. Particular concerns of ours are whether a centralized resource would be equitably available to all AI researchers in a minimally bureaucratic manner. AI is a nascent field but it is also computationally intensive. It is therefore important that there is plenty of scope for exploratory types of research for which one cannot write the typical applications that are required for NSF XSEDE or DOE INCITE computing allocations. By its very nature, one doesn't know if these exploratory studies will lead to useful results, but it is only by exploring that advances will be made.

Some more specific responses to the listed questions are below.

1. *What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]*

A) A NAIRR should aim to provide resources to all federally funded researchers in as simple, transparent, and equitable a manner as possible. Metrics should include delivered computing cycles as well as equity of access.

C) Funding agencies supporting the area should not separate research funding (e.g. NSF grants to support students) from the computational resources needed to undertake that research. Funding agencies should also be mindful of the exploratory, but computationally intense, nature of a lot of AI research and allow for a much more flexible allocation process than exists at current national computing facilities which are driven by proposals with formal milestones.

D) A central priority should be compute resources that are easy to access. There are already many avenues for education and training in AI that can be leveraged.

2. *Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?*

The amount of computing that can be delivered and the availability of a variety of hardware (e.g. a balanced combination of CPUs, GPUs, TPUs, FPGAs) are important.

3. *How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?*

A critical point regarding equity and diversity is to note that a peer-reviewed, proposal-driven system will favor established researchers in the field. If there is a proposal-driven process, it

should have a “junior researcher” track (with a defined fraction of computing available through it, not resulting in average allocations that are far smaller than regular allocations). Requests for computing and eventual allocations should be monitored carefully with regard to diversity.

4. *What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?*

N/A

5. *What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?*

If public-private partnerships provide a cost effective means for stable and flexible resources, they should be encouraged.

6. *Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?*

Requiring additional proposals for science that has already been peer reviewed and funded is a limitation and is not democratic. A national facility, if set up in the right way, could help democratize access by facilitating exploratory AI research in directions that the community has already identified as having strong scientific potential.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

The MITRE Corporation

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Introduction

As a not-for-profit organization, The MITRE Corporation works in the public interest to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation through the operation of multiple federally funded research and development centers and labs, and participation in public-private partnerships. Working across federal, state, and local governments—as well as industry and academia—gives MITRE a unique vantage point. MITRE works in the public interest to discover new possibilities, create unexpected opportunities, and lead by pioneering together for public good to bring innovative ideas into existence in areas such as artificial intelligence (AI), intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience.

MITRE has a long history of partnering with federal agencies to apply the best elements of AI and machine learning (ML) while developing and supporting ethical guardrails to protect people and their personal data. Our team is committed to anticipating and solving future needs that are vital to the success and safety of the public and the country. In crafting this response, MITRE has drawn upon its rich experience in leading Government technology and policy solutions in healthcare, cybersecurity, AI assurance, social justice, identity management, systems engineering, and being an “honest broker” for data.

In the following pages, we offer thoughts on the National Artificial Intelligence Research Resource (NAIRR) roadmap elements (Question 1), including opportunities and goals (Element A), infrastructure (Element D), data (Element E), and security (Element F). We also provide input on NAIRR capabilities and services (Question 2), ethical considerations (Question 3), and democratization (Question 6). MITRE is eager to engage further with the NAIRR Task Force and we stand by ready to support this effort.

Response to Question 1

What options should the Task Force consider for any of roadmap elements A through I ... and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

Element A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success

Recommendation – Extending beyond support to individual AI researchers, the NAIRR should aim to achieve: 1) community-based research collaboration, 2) hypothesis generation, and 3) a research marketplace.

MITRE has relevant experience stemming from developing infrastructure that supports the Million Veteran program for the Veterans Health Administration (VHA).¹ Like the NAIRR, the VHA infrastructure was designed to democratize access to data, tools, and computing resources

to accelerate research. Although the two efforts aim to support different domains (AI vs. genomics), they aim to provide similar capabilities: provisioning and governing access to curated datasets, maintaining privacy, and making computing resources available. Based on the insights we gained while helping the VHA, we identify the following opportunities for the NAIRR.

Community-based Research Collaboration. Solving the wicked research problems of the next century will require community-based approaches. Many researchers will initially be interested in getting access to the NAIRR’s resources so that they can train models germane to their individual research interests. However, the true potential of the NAIRR is more likely to be associated with designing mechanisms and incentives for catalyzing transdisciplinary research collaboration across NAIRR research activities. A community of researchers can be more effective in sharing data-curation methods and can model development approaches, positive as well as negative results, and even real-world performance information of AI models developed with NAIRR resources. There are several models of effective research collaboration that can be used to structure interaction mechanisms among researchers within the NAIRR—for example, see Trochim et al. (2008).²

Hypothesis Generation. As researchers use the NAIRR to tackle scientific and technological challenges that fall on similar lines of inquiry, there will be a significant opportunity to cumulate findings and establish a knowledge base that expands over time. Such cumulation may be automated to a certain degree if the computing environment can be instrumented to identify and extract actionable insights as models are trained. The resulting knowledge base can then be used to build intelligence to suggest what may be worthwhile to investigate next for a given line of inquiry. AI is being used for scientific discovery in many domains (e.g., see Daniels et al., 2021).³ Developing such intelligence can be a fundamental research thrust that the NAIRR actively pursues.

Research Marketplace. Although the NAIRR should provide the necessary incentives for the open sharing and cumulation of data and resources, researchers may want to keep certain types of artifacts private—mainly due to intellectual property considerations. A “research marketplace” could be experimented with where researchers exchange algorithms, models, data, and perhaps even their own services through protocols they control. Emerging bartering approaches in collaborative networks (see Dalli et al., 2019)⁴ may be of relevance to designing such a marketplace.

Element D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure

Recommendation – As a national resource supporting open AI research, it is vital for the NAIRR’s infrastructure and shared resources to be secured on the network. The NAIRR should also promote research at the intersection of AI and cybersecurity – in particular, AI network defense.

MITRE has led several successful efforts over the years producing tools, datasets, and other resources aimed at enabling researchers and cybersecurity professionals to develop and evaluate network defense solutions that will be relevant to the NAIRR—see MITRE ATT&CK and related activities.⁵ These efforts have demonstrated the value and efficiency of using common frameworks and concepts to foster collaboration and rapid progress for cybersecurity and network defense.

The NAIRR also affords the opportunity to advance research at the intersection of AI and cybersecurity. Our experience has taught us that designing, implementing, and maintaining network emulation tools and simulation packages requires upfront investment that can pose a barrier to researchers. The NAIRR could be resourced to provide the shared infrastructure and tools to a) emulate computer networks for the purpose of training AI network defenses at scale and b) simulate large collections of adversarial scenarios for domains such as autonomous vehicles—with the purpose of assessing mitigation of AI-targeted cyberattacks.

Element E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource

Recommendation – The NAIRR should incorporate the following pathways to trusted data sharing: 1) creating privacy-preserving derivatives, 2) making data resources interactive, and 3) setting up data enclaves within the NAIRR. The NAIRR will greatly increase its relevance and impact by incorporating these pathways, and to do so will require the enlistment of a trusted data broker with the experience, reputation, and unbiased independence.

One of the main objectives of the NAIRR is to provide access to large-scale, robust data sets conducive to a broad range of AI research. MITRE recommends the NAIRR incorporate pathways to trusted data sharing. In particular, the NAIRR should enable open research that leverages select sources of sensitive data in a manner that is responsible, reliable, and privacy preserving; thus, augmenting the NAIRR’s data repository of public AI research data sets (which may be limited in number, size, and quality), with more operationally relevant data sets.

The Aviation Safety Information Analysis and Sharing (ASIAS) Program has worked through similar issues.⁶ ASIAS is a partnership between the Federal Aviation Administration, MITRE, and stakeholders across the aviation industry. As a not-for-profit organization, MITRE serves as an “honest broker”—a trusted third-party role that fosters sharing of data amongst industry competitors, which would normally be difficult due to trust issues. The ASIAS model has resulted in a demonstrated ability to protect and appropriately handle sensitive data—a key factor in bringing new operators into ASIAS as the program expands and increases aviation safety for all stakeholders. Based on this experience and track record, MITRE recommends the NAIRR implement the following data pathways.

Creating Privacy-Preserving Derivatives. The NAIRR should selectively incorporate data from sensitive sources, containing sensitive information,⁷ and through techniques such as de-

identification and aggregation, create data set derivatives still useful for tackling open research problems while being privacy-preserving with sensitive data coded or redacted. MITRE has pioneered the use of cryptographic hashing to create non-reversible de-identification of sensitive fields in aviation data while retaining the ability to fuse multiple data sources for more context-aware analysis. For over a decade, MITRE has provided metrics, benchmarks, and analysis products to ASIAs stakeholders while ensuring that competitors' data remains de-identified, without tracing data back to sources.

Making Data Resources Interactive. The NAIRR should give researchers the ability to interact with (query, interrogate, and visualize) data within the NAIRR repository. Experience from the ASIAs Program shows that with careful construction of functions, data interactions can be successfully, responsibly performed directly on sensitive data. ASIAs provides interactive dashboards tailored to stakeholders to meet the needs of authorized users. Each capability has its own infrastructure for preparing the underlying data—some create aggregated extracts of the data to be shown, while others are based on data aggregation performed at the time of rendering. Each has its own security configurations to control data access.

Setting Up Data Enclaves. When the characteristics of a type of sensitive data are not conducive to de-identification or aggregation, then the NAIRR should incorporate the ability to create secured data enclaves that facilitate research by granting access to a controlled set of permitted users. MITRE has a proven track record of advancing the state of the art for solutions in automating the creation of secure analytical enclaves.⁸ The NAIRR will benefit from such capabilities.

Element F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls

Recommendation – As a national resource supporting open AI research, the NAIRR should be resilient and protected against adversarial threats. The Adversarial Threat Landscape for AI Systems (MITRE ATLAS) knowledge base, tool set, and community can be of essential help.

The NAIRR must be cybersecure as a high-valued, heavily invested, networked, national resource. The security challenges are compounded by the way AI systems represent a unique and rapidly expanding attack surface with associated risks not addressed by traditional cybersecurity controls. These attacks can be a highly effective means of exfiltrating sensitive data, stealing intellectual property, or otherwise subverting the AI system (and in this case, the NAIRR) for malicious purposes, even when the AI models meet traditional assurance standards and are effectively secured from cyberattacks.

Consideration for adversarial attacks on the NAIRR should include issues such as data poisoning of hosted data sets, propagating unknown trojan horses and vulnerabilities within AI models and tools, and fostering research contributing to the development of technologies for nefarious applications. The source of threats to the NAIRR is also diverse, ranging from nation-state actors to cyberterrorists to overly adventurous students. NAIRR administrators must be proactive and not wait for a major crisis before they begin addressing such threats.

MITRE, utilizing input from Microsoft and a broad coalition of private sector companies, developed the Adversarial Threat Landscape for AI Systems (MITRE ATLAS).⁹ ATLAS was developed by synthesizing real world case studies, voluntarily submitted by a wide range of industry partners, which detail real attacks conducted against AI systems. The ATLAS team used this to build a robust, common taxonomy of attack tactics and techniques that map to a broad range of contexts to empower security analysts across industry and within the government to detect, respond to, and remediate threats against AI systems.

Since its release in the Fall of 2020, ATLAS has rapidly impacted AI security across multiple industry verticals, with industry teams such as Microsoft, Bosch, Ant Financial Group, and Airbus testing their own AI systems using the ATLAS model. After using the model, these companies contributed the results of their tests as case studies to further improve the ATLAS knowledge base. Additionally, this collaboration resulted in Microsoft's release of a powerful open-source tool set called "CounterFit"¹⁰ based on ATLAS, which gives companies and organizations who cannot afford dedicated AI security practitioners a robust ability to evaluate their own AI-enabled systems against known AI threats. The NAIRR should take full advantage of the growing ATLAS knowledge base and tool set, and the Task Force is encouraged to have administrators of the NAIRR participate and contribute to the ATLAS community.

Recommendation – The NAIRR should utilize federated identity technologies for its management of access controls.

MITRE has been working on the emerging trend of *decentralized technologies*¹¹—in particular, *federated identity*¹²—which is the current direction for Identity & Access Management.¹³ This is an ideal option for open platforms such as the NAIRR. Federated identity refers to linking an individual's electronic identity, authentication, and personal attributes across multiple web service endpoints using Single Sign-on. The website being accessed trusts the identity provider to validate their credentials using industry standard security protocols such as SAML 2.0, OpenID Connect, and OAuth 2.0. The current work MITRE is engaging in is related to the ID.me platform, which uses a federated identity model for the Login.gov identity and authentication service offering.

Response to Question 2

Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Recommendation – The NAIRR should provide the following data and privacy-preserving capabilities—synthetic data generation and federated learning.

The NAIRR can benefit from lessons learned from MITRE's work in facilitating research on health data. Health data is very tricky. Legally, it cannot be used for non-treatment purposes without patient consent with limited exceptions under carefully controlled circumstances.¹⁴ If

patient data is shared, subsequently exposed, the potential penalties are severe, and include criminal penalties. Health systems that possess patient data are very hesitant to share it, making it unlikely that a government agency or contractor will be able to amass health data for study purposes. However, there are ways to make sensitive data like health data more accessible, and the NAIRR should provide these as services and capabilities.¹⁵

Synthetic Data Generation. In addition to hosting data sets, the NAIRR should provide tools so researchers can augment their experimentation by generating their own synthetic data at scale. MITRE has extensive expertise in synthetic data generation. Its Synthea tool “is an open-source, synthetic patient generator that models the medical history of synthetic patients,” providing “high-quality, synthetic, realistic, but not real, patient data and associated health records covering every aspect of healthcare. The resulting data is free from cost, privacy, and security restrictions, enabling research with Health IT data that is otherwise legally or practically unavailable.”¹⁶ With the NAIRR providing synthetic data generating services such as Synthea, researchers will have access to a scalable source of freely, unencumbered data; however, there are trade-offs to using synthetic data over real data. Synthetic data tends to be cleaner, more homogenous, and more structured than real data. Neural networks trained on synthetic data in the lab are at risk of not performing as well on real-world live data. Nonetheless, NAIRR researchers would benefit from on-demand supplies of scalable and tailorable synthetic data.¹⁷

Federated Learning. The NAIRR should support federated learning as an alternative to only training AI models monolithically on a centralized repository of data—particularly when facilitating research on more sensitive types of data like health data.¹⁸ Federated learning distributes the AI model training such that the training data stays behind firewalls, providing greater security controls for data owners and stewards. It does require the different data holdings to use a standardized data model. This is an approach used by Observational Health Data Sciences and Informatics Community and several similar research networks.¹⁹

Recommendation – The NAIRR should host community challenge problems of national importance. For example, orchestrating solutions using AI to protect critical infrastructure based on the NAIRR providing a robust modeling and simulation (M&S) engine.

Protecting critical infrastructure—particularly public water infrastructure—is a national priority for the White House and the Hill. It is called out in the \$1 trillion infrastructure bill and the Drinking Water and Wastewater Infrastructure Act. MITRE is actively working to support Government sponsors in the application of AI to protect critical infrastructure with an initial focus on water treatment plants.²⁰ This involves creating a computing infrastructure that supports collaborative experience and access to real-world modeling and simulation (M&S) objects to augment human intelligence and decision making through applied AI.

MITRE recommends the NAIRR provide M&S of critical infrastructure systems that facilitate the broad testing, evaluation, and development of AI-driven analysis, insight, and learning. Critical infrastructure M&S would encompass key aspects of urban environments where water, food, energy, health, finance, security, transportation, communication, and natural systems converge into complex interrelationships and communities. Many of these M&S objects

communicate through a growing internet of things network, where AI-deployed analytics are becoming increasingly important for automation and augmentation of human decision making.

The goal is to generate realistic environments where monitoring, management, security, operational efficiencies, planning, and various other AI augmentation could be creatively and efficiently developed. For AI augmentation to be successful, it is vital for the NAIRR to facilitate and provide sufficient amounts of detailed, labeled data for training AI critical infrastructure approaches. Providing sufficient computing and M&S access will enable underserved AI critical infrastructure collaboration, analysis, and deployment. Further, MITRE recommends the NAIRR prioritize a standards-based, open-federation that emphasizes the joint use of M&S and AI pattern-detection approaches to enable the utilization of public and private digital resources.

Response to Question 3

How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Recommendation – The Task Force should consider establishing an AI review board (AIRB) process for the NAIRR, similar to Institutional Review Boards (IRBs), to promote equity and fairness. The NAIRR should promote research of mechanisms for accountability, transparency, and operational monitoring of AI technologies.

The challenges when developing AI are diverse and can be both technical and social in nature. As a result, no one person or discipline can singlehandedly “un-bias” AI or make AI ethical. The problem is especially relevant in AI where design teams tend to share many attributes (i.e., similar education and degrees, life experiences, and cultural backgrounds). If AI developers do not actively work to incorporate other valid perspectives into the development process, we risk having the AI reflect limited perspectives about how the technology will be used and by whom.

One potential avenue to promote diverse stakeholder involvement in AI research and development (R&D) within the NAIRR is to establish an AI review board (AIRB) that functions similarly to Institutional Review Boards (IRB) for human subjects’ research.²¹ Such review/assessment should also consider baseline criteria for acceptable performance and risk. Baseline performance criteria should be both mathematical and contextual, and criteria should include the perspectives of all affected stakeholders. Risk assessment criteria should guide decisions about the AI’s suitability to a given application domain or intended use, including the level of clarity that different stakeholders require; risk criteria specific to those groups with greater needs; and guidance for higher-stakes cases when legality, ethics, or potential impact are of concern. These AIRBs should be composed of ethicists, subject matter experts, and representatives from communities who will be affected by the deployed model. In addition to “go/no go” research decisions, the AIRB should specify checks and balances with safeguards throughout the lifecycle of the AI technology to identify and mitigate biases/risks. Reviews will need to include data sheets,²² AI model cards,²³ and impact assessments.

The NAIRR should also be used to conduct research into managing biases and risks with AI technologies by studying mechanisms and establishing best practices for accountability, transparency, and responsibly monitoring the impact of AI technologies once deployed in operational use.

Recommendation – The NAIRR Task Force should develop partnerships with industrial organizations that develop massive AI models affording AI researchers access necessary to study issues of equity, fairness, bias, accountability, etc.

Some of the most significant achievements in AI over the last few years have required massively large investments beyond the reach of most research institutions. For example, GPT-3, a language model developed by OpenAI, has 175 billion parameters, and required several thousand petaflop/s-days (10^{15} neural net operations per second for one day) of computation during training.²⁴ Scholars have expressed multiple concerns about the size and opacity of such models.²⁵ One concern, relevant to the NAIRR's goals, is that access to infrastructure to develop, apply, and evaluate such models is currently limited. For example, access to systems enabled by GPT-3 is currently moderated by OpenAI and GitHub who can withdraw access without public scrutiny or recourse.²⁶ This is an issue because researchers are not currently able to sufficiently research the limitations of these one-of-a-kind models nor sufficiently study the implications of their use. Specifically, researchers should have the ability to develop an understanding about how to prevent the misuse of such technologies, for example, by resourceful actors. Partnerships are recommended between the NAIRR and industrial organizations (that develop these massively opaque AI models) such that AI researchers may have the access necessary to study issues of equity, fairness, bias, accountability, etc. The Task Force should consider developing initiatives that foster access for researchers under the NAIRR umbrella to these large industry models and related technologies.

Response to Question 6

Where do you see limitations in the ability of the NAIRR (National AI Research Resource) to democratize access to AI R&D? And how could these limitations be overcome?

Recommendation – the Task Force can accelerate democratization of AI R&D by leveraging the NAIRR to: 1) overcome a lack of educator training in AI and shared educational resources; 2) provide AI training content that is interactive, easy to find, modify, and share; 3) supply access to scalable cloud-based computing, equitable training data sets, and academic journals, all easily accessible over the Internet; and 4) promote an educational culture with greater risk tolerance and flexibility, creating structures and incentives for more educators to introduce AI into their curricula and classrooms.

Democratizing access to AI research and development must include equitable access and awareness, training and support, and adequate resources within the NAIRR. For people to gain

benefit from access to this infrastructure, they must first be provided equitable access and exposure to knowledge about AI itself and to realize its benefits and application. MITRE believes this exposure should occur during the formative education years and independent of one's field of study.²⁷ We believe this vision is limited in at least four areas:²⁸

Lack of training. Lack of training on AI concepts and applications can lead to hesitation and lower confidence among educators to adopt AI material in the classroom. To overcome the tentativeness of teaching unknown material, educators across disciplines must be engaged directly and be provided with foundational training targeted towards specific competencies. Additionally, educators acting in isolation are limited by a lack of support and content sharing. The NAIRR's shared research infrastructure should encourage community support and content-sharing across all levels of the U.S. education system. Specifically, the NAIRR can help alleviate barriers to communication and support by providing a platform that encourages inter-institutional data and code sharing and promotes streamlined communication between learners, educators, and researchers.

Quality educational content. Currently, there is a lack of freely available AI/ML content that is interactive, easy to modify, and appropriately challenging. This means that content should be accompanied by supplementary material suggestions and lesson guidance for educators, allowing them to scaffold lessons by altering lesson vocabulary, contextual reading and discussion questions, and assignment difficulty as needed. Educators are also time constrained in lesson research, lesson planning, and mapping new lesson content to their curricula. Thus, the NAIRR's shared research infrastructure should serve as an environment in which educators and researchers can quickly research, obtain, modify, create, and share existing lessons that meet these criteria. The NAIRR can provide the capability to freely share modified material and additional lesson guidance. This will reduce the time and expertise needed for quality educational content, tailored to the subject area of interest, while encouraging community-based learning models.

Technical resources available. Students, educators, and researchers are limited by access to technical resources including analytic environments, high quality data, and required technology and hardware. For instance, socioeconomically disadvantaged students may be limited to engaging with content available only on a mobile phone through a free cloud-based analytics platform, while other institutions may conduct research using costly analytics software and leverage access to academic journal publications requiring paid subscriptions or other resources. These limitations could be alleviated by providing an analytics infrastructure that is web-based, accessible on mobile phones, and integrates with, or is centered around free or open-source tools and training data. Additionally, the equitability of training data should be considered in lesson development and research, especially in natural language processing applications, as training data sets are often derived from inequitable sources and can result in the perpetuation of biases.

Risk tolerance and flexibility. Educators have uneven levels of systematic incentives and flexibility that allow for them to put in the necessary time and energy to learn, use, and to build AI educational lessons for students. There is an abundance of interest in widening opportunities for students in data science and AI capabilities, but without the necessary top cover, flexibility in curriculum development, easing of existing time demands, and the establishment of personal growth opportunities and recognition, educators cannot easily participate in programs that might

otherwise allow for them to introduce AI into their learning environments. Creating these structures, incentives, and flexibility for educators will further democratize access to AI by enabling more educators to participate in its development and offer educational opportunities to more students.

¹ Million Veteran program. <https://www.research.va.gov/mvp/>. Accessed September 27, 2021.

² Trochim, W. M., Marcus, S. E., Mâsse, L. C., Moser, R. P., & Weld, P. C. (2008). The evaluation of large research initiatives: A participatory integrative mixed-methods approach. *American Journal of Evaluation*, 29(1), 8–28. doi:10.1177/1098214007309280

³ Daniels, M., Toney, A., Flagg, M., and Yang, C. (2021). Machine Intelligence for Scientific Discovery and Engineering Invention. Center for Security and Emerging Technology. <https://doi.org/10.51593/20200099>

⁴ Dalli, D., & Fortezza, F. (2019). The new face of bartering in collaborative networks: The case of Italy’s most popular bartering website. In *Handbook of the Sharing Economy*. Edward Elgar Publishing.

⁵ MITRE ATT&CK is a globally accessible knowledge base of adversary tactics and techniques based on real-world observations; see <https://attack.mitre.org/>.

⁶ ASIAS Stakeholders. <https://portal.asias.aero/>. Accessed September 27, 2021.

⁷ Examples of sensitive information are personally identifiable information (PII) and protected health information (PHI).

⁸ See MITRE Patent No. 10,795,709 – Systems and Method for Deploying, Securing, and Maintaining Computer-Based Analytic Environments.

⁹ Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS). <https://atlas.mitre.org/>. Accessed September 27, 2021.

¹⁰ Pearce, W., Shankar, R. and Kumar, S. AI security risk assessment using Counterfit. May 2021. <https://www.microsoft.com/security/blog/2021/05/03/ai-security-risk-assessment-using-counterfit/>. Accessed September 27, 2021.

¹¹ Anderson, M. Exploring Decentralization: Blockchain Technology and Complex Coordination. <https://jods.mitpress.mit.edu/pub/7vxentm3/release/2>. Accessed September 27, 2021.

¹² How Federated Identity Can Increase access to Services and Benefits Online. ID.me. <https://insights.id.me/wp-content/uploads/2020/06/How-Federated-Identity-Can-Increase-Access-to-Services-and-Benefits-Online-1.pdf>

¹³ IdAM in a Nutshell. <https://public.cyber.mil/idam/idam-in-a-nutshell/>. <https://public.cyber.mil/idam/idam-in-a-nutshell/>.

¹⁴ Limited Data Set (LDS) Files. <https://www.cms.gov/Research-Statistics-Data-and-Systems/Files-for-Order/LimitedDataSets> Accessed September 27, 2021.

¹⁵ Another benefit of synthetic data generation is to extrapolate data values that are not routinely collected or that could adversely impact individuals if asked, such as religious affiliation or LGBTQ+ status. The NAIRR can also benefit from MITRE’s work with the “All of Us” research program on ensuring representative data. (See <https://allofus.nih.gov/>)

¹⁶ Synthea Empowers Data-Driven Health IT. <https://synthetichealth.github.io/synthea/#about-landing>. Accessed September 27, 2021.

¹⁷ There are other options besides Synthea which you can find by searching on “synthetic health data.”

¹⁸ M.J. Sheller et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci Rep* 10, 12598 (2020). <https://doi.org/10.1038/s41598-020-69250-1>

¹⁹ OHDSI. <https://ohdsi.org/>. Accessed September 27, 2021.

²⁰ MITRE’s own work in this area is called MITRECity.

²¹ For Office for Human Research Protections’ *IRB Guidebook* see http://wayback.archive-it.org/org-745/20150930181805/http://www.hhs.gov/ohrp/archive/irb/irb_guidebook.htm.

²² For example, see <https://www.microsoft.com/en-us/research/project/datasheets-for-datasets/>.

²³ For example, see <https://modelcards.withgoogle.com/about>.

²⁴ T. B. Brown et al., *Language Models are Few-Shot Learners*. 2020.

²⁵ See for example, E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021, pp. 610–623.

²⁶ Quach, K. A Developer built an AI chatbot using GPT-3 that helped a man speak again to his late fiancée. OpenAI shut it down. September 2021. https://www.theregister.com/2021/09/08/project_december_openai_gpt_3/. Accessed September 27, 2021.

²⁷ To this end, MITRE established Generation AI (Gen AI)—a consortium of faculty and students—preparing educators today (and tomorrow’s workforce) to be creative, ethical problem solvers competent in the application of AI technologies. Gen AI is empowering teachers across the country through sharing lesson plans, curated data, and other classroom materials.

²⁸ These recommendations have been informed by MITRE’s work on its Social Justice Platform, which “provides resources, data, tools, and frameworks that empower decision makers to create and sustain equitable solutions that bring positive change for a more just society.” See <https://sjp.mitre.org/>.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Moffitt Cancer Center

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

09/30/2021

Subject: RFI Response: National AI Research Resource

We read with interest the RFI of the National Artificial Intelligence Research Resource (NAIRR) Task Force to develop a roadmap for a nationally shared research infrastructure. This initiative is expected to provide Artificial Intelligence (AI) and Machine Learning (ML) researchers and trainees across the scientific disciplines with access to computational resources, high-quality data, educational tools, and user support. We at Moffitt Cancer Center full-heartedly support such an initiative, as these resources are much needed to advance AI/ML in general, as well to tackle the chronic pandemic of cancer through proper application of this powerful technology.

On one hand, AI/ML has been presented as a savior technology that will transform oncology practice by automating laborious routine tasks, improving efficiencies, reducing costs, and enhancing decision-making support of complex oncology processes, including everything from treatment management to prediction of outcomes and adapting prescriptions over the course of treatment. On the other hand, the anticipated implementation of AI/ML in oncology and healthcare in general has been limited in scope and sometimes stagnant with less than 5% of major healthcare providers implementing any form of AI/ML solutions (HIMSS). Moreover, fairness and ethical concerns have casted shadows on the ability of AI/ML to reduce health disparities versus exacerbate existing inequities. Recognizing the potentials of AI/ML to transform cancer care and the need to address barriers in the realization of its potential, we have established the first AI/ML department in oncology at Moffitt Cancer Center, which was featured in the Cancer Letter ([Vol.47 No.07, February 19, 2021](#)). A key mandate of establishing this department is the integration with the parent cancer center scholarly activities in research and the facilitation of AI/ML clinical translation as part of the center's data/digital ecosystem, by enabling AI/ML driven multi-scale convergence of patient data into actionable cancer research and daily oncology care.

Given the aforementioned background, we provide our responses to the questions posed by the NAIRR's RFI, which we hope would be helpful in developing the proposed roadmap while highlighting specific cancer needs:

1. *What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]*
 - A) The vision presented by the task group and the lessons learned for a secure single access platform and friendly presentation are welcomed. However, additional lessons learned from the TCIA/TCGA repository are that proper data structures and proper annotations for AI/ML are needed for successful utilization. In addition, indicators of data quality may be required when such data are used for training or testing of existing and/or new AI/ML algorithms and technologies.
 - B) Data governance and open access have been a challenge in the past, even through federally supported entities, including a multi-institutional cooperative groups, where democratization of data is [part of their mission, were unfortunately mired with regulatory

and technical challenges. Hence, it would be useful to have clearer and transparent criteria for resource sharing and data access eligibility and an independent oversight committee to organize such processes in a fair and transparent manner.

- C) The governance model needs to be representative of stakeholders of the different disciplines including sensitive sectors such as healthcare. We advise that an organizational structure should be created that can provide equal opportunity access given the expected high demand. This can be done in a similar fashion to how federal granting agencies currently provide resources. For instance, the application process can be structured in a hierarchal manner from the participating institutions through a nomination process to ensure that resources are properly allocated, and the use cases will have the expected impact on the targeted application area. Additionally, priorities can be provided within this process to meet national interests for AI/ML applications in the different sectors.
- D) The compute capability for developing and applying AI/ML already exists in the cloud-based platforms noted by the RFI. However, a major hinderance is the for-profit nature of many of these platforms; subsidization may be needed for academic and research use. As in response to A), curated data sharing is another major challenge; support for federated learning infrastructures may be necessary in the long run to enrich the available data resources while respecting legal and ethical concerns.
- E) Making high-quality, government funded datasets available is an important objective. However, there are several lingering challenges to this objective, which include the availability and open access due to privacy concerns and other regulation and technical issues. Supporting institutional infrastructures for secure federated data access may offer a feasible solution and facilitate resource sharing in an equitable manner.
- F) Proper cybersecurity safeguards would need to be evaluated and necessary requirements should be presented in a clear and transparent matter. This will provide the opportunity to replace antiquated security protocols with more modern approaches that can protect the data from unauthorized use without significantly hindering access and utilization.
- G) Privacy in healthcare is a major concern and standardization of protocols for anonymization and harmonization of the process would be an advantage for a centralized setting. In the federated setting, anonymization may not be necessary if proper security guards are put in place that can protect the data from unauthorized use while facilitating secure access.
- H) It is important to have sustainable infrastructure with long term objectives for hardware and software upgrades as the nature of AI/ML is to evolve over time and the need for better computing hardware and software are likely to persist and increase.
- I) The required parameters for sustainability need to meet the needs of the different sectors and their pressing use cases. This is especially true in the domains of healthcare and cancer research, where such needs continue to evolve over time. Moreover, a repository to maintaining codes, such as GitHub, with better security guards may be needed as well for such sharable resource too.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Provision of curated datasets has been a major limitation for widescale implementation and validation of existing and newly developed AI/ML tools at academic centers. ImageNet has been a driving force for the latest AI success, and something similar would be essential for benchmarking different algorithms across the different sectors especially in biomedical sciences. If too many resources are required for building an exhaustive AI/ML database, subsets of government funded and curated datasets could be sequestered for evaluations purposes only, which would still be very valuable for having common benchmarking schemes across the board for newly acclaimed AI/ML algorithms that can assess their robustness in a systematic fashion.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

There are tools, though still evolving, for evaluating AI/ML fairness that need to be encouraged and utilized as part of the AI/ML lifecycle. These tools include the interpretability and explainability evaluation tools. Moreover, the datasets provision of item D) can also be utilized as a testing bed of the fairness of the developed AI tools and how representative of the society and its underrepresented minorities. These are likely to be use case dependent but general principles can be chartered as suggested by the recent European Union and the FDA guidelines for trustworthy AI. In addition, specific datasets that include underrepresented populations may need to be generated as part of the implementation process for human biomedical research such that AI/ML algorithms trained/tested on those data are more broadly generalizable. Indeed, optimizing the generalizability of AI/ML to more diverse populations can accelerate the narrowing of the current disparity gap instead of exacerbating it as has been postulated

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

There are extensive compute resources, such as the ones presented by the STRIDE or the XSEDE supercomputer platforms, that can be leveraged and expanded to include more institutions. The main current limitation is in the availability of proper datasets for training and validation across the various domains, especially in the case of cancer research, which may limit the options for standardized vetting of newly proposed AI/ML techniques and subsequent translation into the clinic and daily care.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

The STRIDE example, if it can be generalized, may be a good one to follow across multiple disciplines. However, affordable chargeback policies may need to be implemented with eth commercial vendors. Moreover, challenges associated with existing data commons such as the TCIA/TCGA repositories including proper data structures and quality metrics may need to be resolved for more effective utilization

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The limitations may not be technical but mostly related to proper access to datasets that would allow standardized evaluation of newly proposed AI/ML technologies along their lifecycle. Moreover, getting buy-in from institutions to make their existing resources available to NAIRR may be another challenge that may need to be handled through proper incentives and value-based propositions that can lead to the shared positive impact on the society and its institutions including cancer care.

On behalf of Moffitt Cancer Center, Tampa, Florida

Issam El Naqa, PhD, Chair of Machine Learning

Dana Rollison, PhD, Associate Cancer Director and Chief Data Officer

Edmondo Robinson, MD, Chief Digital Officer

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

NASA

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

NASA Feedback RE: National Artificial Intelligence Research Resource (NAIRR) Request for Information (RFI).

I. Strategic Level Inputs:

- NASA sees AI as a powerful capability for the United States, and therefore encourages the NAIRR as it seeks to bolster AI research and partnerships
- NASA's resources are constrained on priority missions directed by Congress and no spare NASA resources are available to contribute to the NAIRR
- NASA flagship missions such as the Artemis campaign rely on technology that has been substantially matured and tested by researchers and technical experts.

II. Detailed Inputs, based on the NAIRR RFI. NASA has embedded answers under each RFI question.

1. What options should the Task Force consider for any of the roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success.

NASA: The options should include establishing:

- *A Center of Excellence to oversee implementation*
- *Training and education program to develop workforce*
- *Framework and processes for reuse and reproducibility of AI experiments*
- *Incentivization for open sharing and reproducible results*
- *A method to rate capability maturity to guide adoption when ready*

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and

NASA:

One possible model is for industry and academia to largely run and resource the NAIRR, with government advocacy, participation, and inspiration

ii. A governance structure for the Research Resource, including oversight and decision-making authorities;

NASA:

Based on B.i. (above), A possible model is for the government to encourage industry and academia to run the NAIRR, to include business & academia providing most oversight and decision-making functions. The government may not have to provide program management,

but could instead focus on NAIRR advocacy, participation, and inspiration. This is one option; not all of NASA's answers to other questions assume this option.

Another option could be the NIST model. See <https://www.nist.gov/artificial-intelligence>.

Regardless of the mechanism for governance, an independent oversight function would be necessary; an organization to ensure the public interest is respected as technologies develop.

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

NASA: The options should include:

- *Long-term funding model to support and manage allocation of resources*
- *Collaborative governance with representation from industry, academia and governments*
- *Strategic planning meetings, reports, success stories, and new challenges*
- *The Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) solicitations framework is an option that could, at least in part, be used a model for governance and managing the allocation of resources*
- *The formation of tightly compartmentalized teams focused on well-scoped and discrete tasks, coming together to discuss and disperse innovations and insights*

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

NASA: Many options to address items in this element are being advanced by industry. We should collaborate with the industry to adapt their solutions to meet our needs recognizing regulatory and legal requirements may require unique investment from the government to realize.

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

NASA Science Mission Directorate's (SMD's) Strategic Data Management Working Group's Report on Groundbreaking Science is an option to refer to which addresses this element.

Also, consider creating policies, procedures and sandbox environments to make it easier for industry, academia, and government to team on shared problems. For example, guidance on sanitization of large sensitive datasets to allow for collaboration outside of government.

NAIRR should iteratively poll the user, participant, and customer community to ensure it is providing what is needed. Deliver what the people want by asking the people what they need.

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls;

As some NAIRR capabilities mature, methods would be needed to on-board them into selected sensitive government programs; mechanisms for traversing sensitivity levels would be necessary.

G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;

H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector;

NASA: NASA's Mission organizations leverage industry and academia via a variety of engagement mechanisms, some of which might allow NASA to become a customer of NAIRR participants. As one example of many, NASA SMD's open source science program includes targeted solicitations, incentives, and collaborations with private sectors.

I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

NASA priority list of capabilities and services are as follows (order is not implied):

- *Curated training datasets and metadata registries*
- *Computing resources*
- *Educational tools and services*
- *Resident expertise, to include AI, ML, scientific domains, cybersecurity, social bias and IT*
- *Scalability of infrastructure*
- *A taxonomy of levels of AI capability*
- *Standards to support interpretability and interoperability*
- *Free open source libraries and APIs for relevant capabilities (*Note: if government investment helps propel the NAIRR, this makes a good case for broad, free sharing)*
- *Secure access control (where necessary)*
- *A user interface portal for accessing and finding resources*

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

NASA Recommends:

- *To develop practical guidance for ethically applying AI and forming mechanisms to encourage & ensure adherence as needed*
- *To develop a checklist for best practices*

- *To conduct independent ethics review of research*
- *To develop systems to examine unintended consequences of AI research*
- *To include additional ethics experts into discussions where decisions are made*
- *Develop a portfolio of federal entities looking to be responsive to actions such as the executive order related to Advancing Racial Equity and Support for Underserved Communities Through the Federal Government, targeting inclusive AI can leverage best practices, industry resources and community contacts for strategic federal mission focused initiatives. Link the NAIRR with other initiatives for complementary effect.*
- *Participate in global and national ethical AI discussion and debate; this area is a work-in progress; NAIRR can both contribute-to and benefit-from the larger ethical AI discussion*
- *The NAIRR and its components should establish a set of standards and guidelines for usage of data sets. These guidelines should define the boundaries for the research using NAIRR, taking issues of ethics and transparency into account. In addition, standards should be established to mitigate impacts of biases inherent within data sets when used by machine learning algorithms*
- *Provide for resident personnel with expertise and knowledge in how AI and ML systems have exhibited bias that has led to social harm – personnel that work integrally with educational tools and services*
- *Recognize that ethics are subjective and vary by culture. Since there is no canonical definition of “ethical”, the goal for any software system should be responsible and transparent behavior. This holds whether that system incorporates AI or not*

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

NASA: The following building blocks exist across the technical community:

- *High end computing platforms for AI*
- *Procurement of cloud computing platforms to scale AI in production*
- *Open data and open source policies*
- *Competitive programs targeted towards AI*
- *Public and private partnerships with relevant partners*
- *The European Union’s coordinated AI plan, which can serve as a benchmark*
- *The Small Business Innovative Research (SBIR) and Small Business Technology Transfer (STTR) programs could provide channels for pursuing some NAIRR work*
- *Federal programs and partnerships, such as the NSF National AI Research Institutes, could serve as existing building blocks for the NAIRR*
- *DARPA’s Explainable AI work could contribute to overall NAIRR approaches*
- *NASA’s Space Apps Challenge is a good example of global collaboration and crowdsourcing in emerging technology spaces, to include AI*

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

NASA:

Most of the innovations around AI are happening in the private sector and academia. The role of the partnership is invaluable in bringing the experts and tools and services to solve problems within the government. The partnership can also provide avenues to fill in necessary AI workforce within the government.

NASA's Earth Science Data Systems (ESDS) Program has established strategic partnerships with public and private companies to help further its data management and data development efforts through non-reimbursable Space Act Agreements. Some examples of our success and outcomes can be found at <https://earthdata.nasa.gov/collaborate/esds-public-private-partnerships>.

NASA SMD and NOAA's NESDIS division have formed a working group to address cloud architecture data access and discovery challenges. The group meets quarterly to share knowledge, and status on joint initiatives focused on developing data governance (provenance) guidelines and common practices through the data life cycle in the cloud. Some examples include a joint cloud-based data expedition pilot to demonstrate interoperability between NASA and NOAA data.

The Frontier Development Lab, an industry-academia NASA-focused external innovation lab, is another good example. The Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) programs is another example that could be used as a model for public-private partnerships in the NAIRR.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

NASA:

Limitations include: seamless access to computing platforms to accelerate AI, availability of high quality training datasets, and reproducibility of AI experiments. To overcome these limitations, we need to prioritize AI across agencies and develop an AI strategy that addresses these limitations. Furthermore, we need better collaborations with academia and the private sector to develop workforce for the future and exchange knowledge. We also need to create an environment where domain experts and AI experts can work together. To accelerate adoption of AI we need to foster the culture of open sharing and collaboration. Some NASA missions include high-cost, low sample size data; techniques for adapting AI to these data sets are needed. The AI community must also find ways to communicate and show the value of AI-enabled solutions to mid-late career workers who are comfortable with more traditional techniques.

Scientific research depends upon collaboration and sharing of results, but heightened security barriers around a system cause obstacles that limit access by the scientific community. While security is critical in certain areas, easily accessible tools and data that can be easily shared to enhance collaboration is a critical aspect to accessible science. One solution to this problem is providing different tiers of access. These would have similar capabilities, but having a "public" tier with public data and compute that is accessible to anyone through standard, community based login methods would make these tools more accessible. There can then be tiers with

more increasing level of security, but with more resources, access to confidential information, and even restricted tools.

Supporting the open science community: Providing publicly accessible data and software under open and permissible license significantly democratizes the access to data. The fact that people anywhere can download the data and software for AI/ML research is probably the greatest democratization process in place. As these are available without any gateways, these are incredibly accessible. However, the teams that are producing the data need the resources to produce data in AI/ML accessible formats with metadata, that is well labeled, and has accessible documentation. Teams that are producing the open source software that underlies much of AI/ML are often not well supported, especially the libraries that are underpinning the basic structure. Sustainable funding to support both maintenance and innovation is necessary for the software development. One further gateway to lower is that the documentation and training is available to make these resources accessible.

Limitations may exist with respect to democratizing access to communities, institutions, and regions that have been traditionally underserved and underrepresented with regard to AI research. Limited access to resources/hardware/infrastructure (i.e, digital divide) necessary to leverage NAIRR capabilities and services may present additional challenges. These limitations could be addressed through targeted outreach to centers, historically-underserved universities, programs, and conferences that place emphasis on engaging underrepresented communities in AI research and educational activities.

Failure to elevate AI as a profession could be a barrier to progress. The NAIRR can bolster National AI progress by fostering AI as an emerging top-tier profession. With academia participation, there is a great opportunity to craft a variety of AI-related professional development materials & activities to foster broad and deep AI knowledge and skills across all elements of the United States. With industry, academia, and government participation in the NAIRR, there is also great opportunity to foster other aspects of AI as a profession, such as ethical AI debate & guidance, fostering powerful & useful standards, self-policing AI-related activities, and hosting workshops / conventions, etc.

Question 6 also generated some discussion regarding whether it is the right question. Should the NAIRR actually focus on democratizing access to AI R&D? Can access to these capabilities even be controlled or limited? Perhaps part of this function would be matching NAIRR resources to customers by expertise, problem set and data type. Another possibility would be for the NAIRR roadmap to focus on overcoming limitations by eliminating economic biases such as those present in our public education system in addition to systemic biases of fielded systems (e.g., Tech-company branded credit cards, facial recognition, COMPAS). Solutions could include exposing researchers, students, and the general public to AI resources such as educational tools and data sets as well as application-appropriate benchmarks for fielded systems.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

National Center for Atmospheric Research

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

RFI Response
**White House Office of Science and Technology Policy and National
Science Foundation**
**Implementation Plan for a National Artificial Intelligence Research
Resource**

From the National Center for Atmospheric Research (NCAR)

Scientists and Software Engineers from NCAR have discussed the RFI on an Implementation Plan for a National Artificial Intelligence Research Resource, and are pleased to have the opportunity to respond.

Operations and Mission of the National Center for Atmospheric Research

NCAR's mission is "to conduct research that contributes to the depth of fundamental understanding of the atmosphere and its interaction with society and the environment and to develop and transfer knowledge and technology that expands the reach of atmospheric science." NCAR has a successful history of transferring technology and knowledge to U.S. government agencies, the private sector, and foreign governments. NCAR is eager to collaborate with other organizations on a shared computing and data infrastructure to provide AI researchers and students with access to a holistic, advanced computing ecosystem.

NCAR is operated by the University Corporation for Atmospheric Research (UCAR), a non-profit organization established in 1960 to oversee a wide range of programs and facilities that support its 120+ university affiliates and the national and international scientific community.

NCAR's current role in advanced computing, data ecosystem, and AI research

NCAR has a long-standing research program in artificial intelligence (AI) and machine learning (ML). That research began under externally funded projects and has helped lead the meteorology and climate community to recognize the efficacy of applying these methods to weather forecasting, hydrometeorology, climate simulation, and various applications, including aviation meteorology, renewable energy, surface transportation,

forecasting in support of agriculture and food security, modeling for wildland fire management, and much more. Much like statistics, AI/ML provides a facility that can be used for understanding data, such as to describe the physics of the Earth system. In addition, these methods enable actionable science, such as by combining physical and social data to infer the likelihood of impacts. While care must be exercised in the training data sets to avoid bias, machine learning is well suited for tasks that are hard to describe formally or too expensive to compute at scale but where data are abundant or easily generated, a realm prevalent in Earth system science. Observation and remote sensing systems are collecting higher spatial- and temporal resolution data, and higher-resolution models are being run with more coupled processes to increase the realism of simulations. Our scientific ambitions overshadow the expected trajectory of our computing resources. Machine learning approaches promise a computationally efficient and scalable means to model the relationships between different data sets with more efficient and scalable computation.

More recently, NCAR has embarked on a major effort to expand the use of AI in environmental systems. For instance, the Computing and Information Systems Laboratory has teamed with other NCAR science labs, including the Research Applications Laboratory, High Altitude Observatory, and Climate and Global Dynamics Laboratory to build, demonstrate, and validate ways to replace certain physical models in our various modeling systems with AI emulators. The preliminary results of these initiatives are encouraging and are expected to lead to significant advances in model accuracy and efficiency.

Q1: What options should the Task Force consider for any of roadmap elements A through I above, and why?

Answers to roadmap topic D

- i. **Develop an ecosystem of interoperable compute and data resources with easy and equitable access** - To provide Artificial Intelligence (AI) researchers and students across scientific disciplines with access to computational resources, access to the current national compute ecosystem needs to be simplified and hurdles that limit equitable access need to be reduced. Despite the efforts of XSEDE, the step of moving from on-campus compute resources to cloud and national resources is still nontrivial. A national effort can reduce this hurdle by providing consistent AI software stacks, user interfaces like Jupyterhub, and the integration and support of cloud approaches. Connecting, leveraging, and expanding the current national infrastructure, primarily NSF-funded resources, together with increased training and support, will democratize access for AI researchers. Domain-specific computational and data resources similar to the

research data archive¹ at NCAR provide the weather and climate research community with connections to other NSF-funded centers integrated by cloud technologies; these could become the building blocks of a national AI compute data infrastructure. NCAR has played a crucial role in public-private partnerships, as shown in answer to question 5. NCAR is well-positioned to play a leading role in connecting the public, private and academic sectors to advance the state-of-the-science in AI research applied to Earth System predictability.

- ii. **Develop shared public datasets and environments for AI training, testing, and benchmarking** – This is undoubtedly happening in the weather and climate community and outside of it. Several organizations interested in advancing the use of AI within the weather and climate community are actively working to identify and deploy high-quality datasets for AI training and comparison to increase the use of AI. For example, the American Meteorological Society (AMS) has fostered a Committee on the Application of Artificial Intelligence in the Environmental Sciences since the late 1990s and began teaching short courses on the topic in 2001, eventually archiving the lectures in a book. In 2008, that committee began holding AI forecasting contests, providing a shared dataset that researchers could use to test their algorithms. More recently, the AMS AI Committee has been banding together with other organizations to provide a complex series of datasets that researchers can explore and compare their applications against peers. When the committee turned to the Kaggle competition website in 2014 to host their contest, they were surprised to find that some non-meteorology AI experts applied techniques that outperformed those currently being applied in the weather community. Those novel techniques were quickly adopted by many others doing parallel research. Thus, the weather community has evidence that such common datasets are effective in helping to advance the state-of-the-science. The data sets will also provide an opportunity for computational scientists and data and compute systems researchers to inform the procurement of the next generation of national HPC systems.

In FY2021, NCAR was selected to serve as the long-term data repository for observations made by the NSF-funded Community Instruments and Facilities (CIF) program. This repository, known as the Geoscience Data Exchange (GDEX), will be an expansion in scope and capacity of the existing DASH Repository. GDEX will provide archival storage, data discovery services, open access, and citable DOIs for CIF data. NCAR is well-positioned to play a leadership role in providing Earth System data sets for the AI research community. Additionally, NCAR has partnered with a commercial cloud vendor to host a subset² of the Community Earth System Model Large Ensemble³ (CESM LENS) dataset together with example notebooks demonstrating how to use the data. This approach democratizes the access to access and use of this data.

- iii. **Expand the capacity of the national compute and data ecosystem to support the demands of AI** - The current NSF national systems are over-allocated and will not be able to provide enough capacity for the demands of

¹ <https://rda.ucar.edu/>

² doi:10.26024/wt24-5j82

³ doi:10.1175/BAMS-D-13-00255.1

the AI research community. Investments in compute capacity as well as in innovative allocation and access models are necessary. An increase of capacity of high-quality scientific themed data repositories connected to advanced HPC and AI cyberinfrastructure will be necessary to support high-resolution output of community model runs like the NCAR-led Community Earth System Model (CESM) that then can be used for the AI research process. Machine Learning meta-data and new data standardized to enable automatic consumption by the AI research process needs to be created as part of the data design process. The expansion of the ecosystem should be driven by a co-design process between the AI research community, the data archives, and the infrastructure engineers.

- iv. **Transdisciplinary collaboration** - The creation of a national AI research resource will require close collaboration among modelers, theorists, experimentalists, engineers, humanists, social scientists, computer and data scientists to create the data sets, training, and compute and data infrastructure necessary for the advancement of AI research. For example, storm-resolving model ensemble runs on pre-exascale and exascale systems will generate massive amounts of data that are too large to store fully. In situ training or lossy compression tailored for AI are approaches that need to be explored as part of the infrastructure roadmap. Convening the Earth System Science community and bringing different communities together is one of the strengths of NCAR, who would be glad to participate in leadership roles in this domain.
- v. **User interfaces** User interface design has produced alternatives to the command line interface such as Jupyter notebooks and science portals (eg. Open OnDemand and Science Gateways) that reduce the learning curve for first-time users, enhance collaboration, and simplify visualization, documentation, dissemination, and reproduction of scientific results. As these tools evolve, researchers need to navigate the perennial tradeoffs between ease of use vs. greater control over the underlying hardware and software. One of the goals of the national AI research resource should be to make AI at scale easy for researchers to conduct.
- vi. **Training, education, and workforce development** - The nation should leverage existing educational tools at different academic and research institutions to create a central repository/resource for ML education and training. Resources should focus on training upcoming scientists and researchers in ML (i.e., students, early-career), as well as training, experienced scientists who are interested in applying ML in their respective areas of expertise. Examples of educational resources can include tutorials and summer schools hosted by universities and research laboratories and educational modules (following the UCAR COMET program). The AI4ESS summer school in 2020⁴ and the TAI4ES summer school in 2021⁵ are examples of educational activities that are necessary to bring participants up-to-speed on how to develop trustworthy AI for the Earth & environmental sciences.
- vii. **Better understand the national AI R&D workforce needs** – AI is becoming critical to our scientific approaches. We need people capable of building and

⁴ <https://www2.cisl.ucar.edu/events/summer-school/ai4ess/2020/artificial-intelligence-earth-system-science-ai4ess-summer-school>

⁵ <https://www2.cisl.ucar.edu/tai4es>

maintaining the computational infrastructure necessary to continue to push the state-of-the-science. These people must know how to deal with “Big Data” at all levels, including how to manage datasets. When AI is used in production involving the types of big datasets that we often use in meteorological applications (e.g., blending in situ and remote observations with output from many models smartly to improve forecasts), this aspect becomes as important as producing the new AI algorithms. This personnel trained in AI and Big Data must be prevalent in public positions as well as in private enterprise. In the meteorology community, several universities are providing training in AI at the undergraduate as well as graduate levels. National laboratories must emphasize applications of AI. At NCAR, we have re-invigorated a cross-laboratory effort in AI, both in basic research and applications that we have fostered over the years. We firmly believe that applying these principles more broadly across a range of topics is needed to assure continued advances in our understanding and in our applied technology.

Q2: Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Prioritization in the following order (see details of the buildings block in the previous answer):

- Priority 1: (iv) Transdisciplinary collaboration: Only close collaboration between a diverse set of domains of expertise will enable a democratized and equitable AI research resource.
- Priority 2: (ii) Develop shared public datasets and environments for AI training, testing, and benchmarking: To advance the use of AI, it is essential to identify and deploy high-quality datasets that are sufficiently described for automatic model training and testing.
- Priority 3: (vi) Training, education, and workforce development: A central training resource is important to provide equitable access to training materials connected to the high-quality data sets to train upcoming scientists and researchers in AI, as well as training experienced scientists who are interested in applying AI in their respective areas of expertise.
- Priority 4: (i) Public and private partnerships: The compute and data ecosystem necessary is only achievable by synergistic partnerships between the private and public sectors.
- Priority 5: (iii) Expand capacity: The current NSF national systems are over-allocated and will not be able to provide enough compute capacity for the demands of the AI research community. Investments in capacity and integration with commercial cloud providers and innovative allocation and access models are necessary. The increase of capacity of high-quality scientific themed data repositories connected to advanced HPC and AI cyberinfrastructure will be necessary to advance AI research.

Priority 6: (v) User interfaces: Evolve the existing JupyterHub ecosystem, to support AI research at scale through an easy-to-use interface.

Q3: How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Our organization has successful examples of approaching this challenge. Specifically, NCAR is a key partner in the NSF AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography (AI2ES). As part of this institute, we are developing novel methods of evaluating the trustworthiness of AI systems for a diverse array of geoscience use cases. Evaluating the trustworthiness of AI systems for environmental science is important because these tools if adopted will be used for high-stakes decision-making that could greatly impact lives and property. A key aspect of the work of this center is involving social scientists to better understand how the stakeholders understand trustworthiness and working directly with end-users to help develop ways to meet their trust. Another critical aspect of this work is clearly defining the components of trustworthiness for a given set of problems and consistently evaluating each system to ensure that trust is warranted. The three focus areas for trustworthy AI in AI2ES are ensuring that AI is explainable, physics-based, and robust.

Weather and climate observation and prediction data are also subject to racial and gender equity biases and AI systems trained with this data could also risk propagating these biases. Globally, weather observations are far denser in the United States and Europe, resulting in poorer characterization of current weather and forecasts in other parts of the world. Satellite measurements have reduced this gap when forecasting larger spatial scales, but there are still systematic biases in observing local weather patterns. Poorer urban and rural areas are less likely to have dense weather observations and may be less likely to report extreme weather events. These data issues could result in AI systems that also perform more poorly in these areas. AI model predictions should be evaluated for these kinds of systematic biases, and targeted observation campaigns may be necessary to collect more verification data in these data-sparse but heavily-populated regions.

Q4: What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Many of these aspects have been mentioned in the answer to Q1. Specifically, some existing building blocks include:

1. XSEDE service providers, such as the NCAR HPC ecosystem that support the atmospheric sciences community.
2. Training efforts at Universities and NCAR. As noted above, many excellent educational resources already exist that could be advertised widely and built upon. Many universities offer on-line classes and prior summer schools such as those offered by NCAR in collaboration with our university members and our partners in the public and private sectors, with archived content for use by future students.
3. Part of NCAR's role in the Earth System Science community is as a convener of the Community - we bring people and resources together to accomplish goals that are important to society.
4. In addition to convening the community, it is essential to provide continuing support. Our own approach in this role for modeling and observational realms is well poised to be scalable to broader communities.
5. As mentioned above, NCAR already provides leadership in the geoscience community for archiving Earth System data. We envision potential to grow services to include preparing digital notebooks that enable accessing the data and assessing the results of ML applications relative to a baseline, i.e., testing and verification frameworks.

Q5: What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Partnerships, including the academic sector in addition to public and private, are imperative to advance NAIRR. Each sector brings important, unique expertise and considerations relevant to their own needs.

As an example, NCAR formerly led a Public-Private-Academic Partnership to Advance Solar Power Forecasting (funded by DOE) that included three other national laboratories, six universities, and 11 private sector partners as well as several international affiliate partners. All sectors participated in crafting the initial vision to ensure that the end product was useful to the end-users and that all research was accomplished in a robust, open, and repeatable manner. This team found that the key to

success was communication and that employing methods from the social sciences facilitated conversation and helped the team members to better understand each other. The final outcome included new models for solar power forecasting, both those based on physical principles and those based on machine learning. It required all sectors working together to adequately design and test their usefulness for the intended application.

Above, we described the current AI2ES partnership with academic, private, and public sectors that is working toward advancing trustworthy ML use in the environmental sciences. In similar ways to these two projects, we believe that the NAIRR will require working across sectors to provide the hardware, software, data, tools, and secure access needed to make NAIRR successful.

Q6: Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The security aspect of NAIRR will be critical to success if the goal of easy and democratic access is to be achieved without compromising software and data integrity. It is difficult to make large systems open to use by students and the public without opening access to bad actors. The right organizations must be involved to ensure the security of necessary data, user information, and applications that are archived.

Secondly, as we wish to democratize AI and data, it will be necessary to train some users in basic computer usage, programming, and database access. The planners should think through how to help users get started without taking an inordinate amount of time from those charged to help users. One opportunity could follow from the user interface perspective highlighted above. Specifically, tools like Jupyter notebooks can lower the barrier to accessing data, algorithms, and compute environments, and these have been successfully demonstrated in training programs for novice users.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

National Energy Technology Laboratory

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Kelly Rose, Madhava Syamlal, Jimmy Thornton, Sydni Credle, and Darren Mollot

In 2020, The U.S. Department of Energy's (DOE) Office of Fossil Energy & Carbon Management (FECM) and the National Energy Technology Laboratory worked with other DOE applied energy stakeholders to evaluate and constrain Artificial Intelligence (AI) research and development (R&D) resource needs for our Applied Energy R&D Office. These findings were used to produce the report that follows, which aligned to the Office of Science and Technology Policy's (OSTP) and the National Science Foundation's (NSF) "Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource." While developed for addressing the AI/ML needs of FECM, the recommendations in the report may address the needs of AI/ML researchers in other areas to support and produce more efficient and effective products and breakthroughs.

Accelerating FECM R&D Using AI and ML

The U.S. Department of Energy's Office of Fossil Energy & Carbon Management (FECM) supported researchers are now applying artificial intelligence (AI) in at least 70 applied research and development (R&D) projects.¹ Most of these projects remain in early development, and all focus on solving known challenges within specific program areas. To evaluate FECM's readiness to better leverage the full power of AI, an internal AI Guidance Team was formed to examine existing gaps in AI capabilities and untapped opportunities for AI to broadly accelerate FECM research.

As a first step, the team reviewed the exceptional AI/ML resources and unique capabilities being developed at FECM's National Energy Technology Laboratory (NETL). Next, the team focused on the hurdles that currently impede or restrict FECM's use of this transformational technology. This paper summarizes the team's findings and recommendations.

Findings of the FECM AI Guidance Team

AI in FE: Hurdles & Recommendation

Data Hurdles. Researchers today spend too much time and effort to find, obtain, and prepare enough high-quality data with the properties and formats needed to drive targeted AI/ML solutions.

Tools Hurdles. It is also hard for researchers to find the right AI tools, so they often develop them from scratch. FE lacks a secure, central hub for developing, storing, and evaluating *models and algorithms*.

Computing Hurdles. Gaining access to the computing power required for AI/ML work can be challenging. Settling for suboptimal computing resources can waste enormous amounts of time, money, and energy.

Recommendation

FE should accelerate efforts to stand up a secure, integrated AI/ML development environment that provides researchers seamless access to the tools, data, and resources required for rapid AI-enabled breakthroughs in the FE mission space.

FECM has no shortage of promising applications for AI, and access to an **integrated AI/ML development environment** could significantly accelerate the pace and broaden the scope overall. A foundational AI/ML development resource that provides vast libraries of curated data; proven, interoperable tools; expert AI support; and access to the necessary computing resources (see Figure 1) will significantly expedite FECM progress and reduce AI project costs.

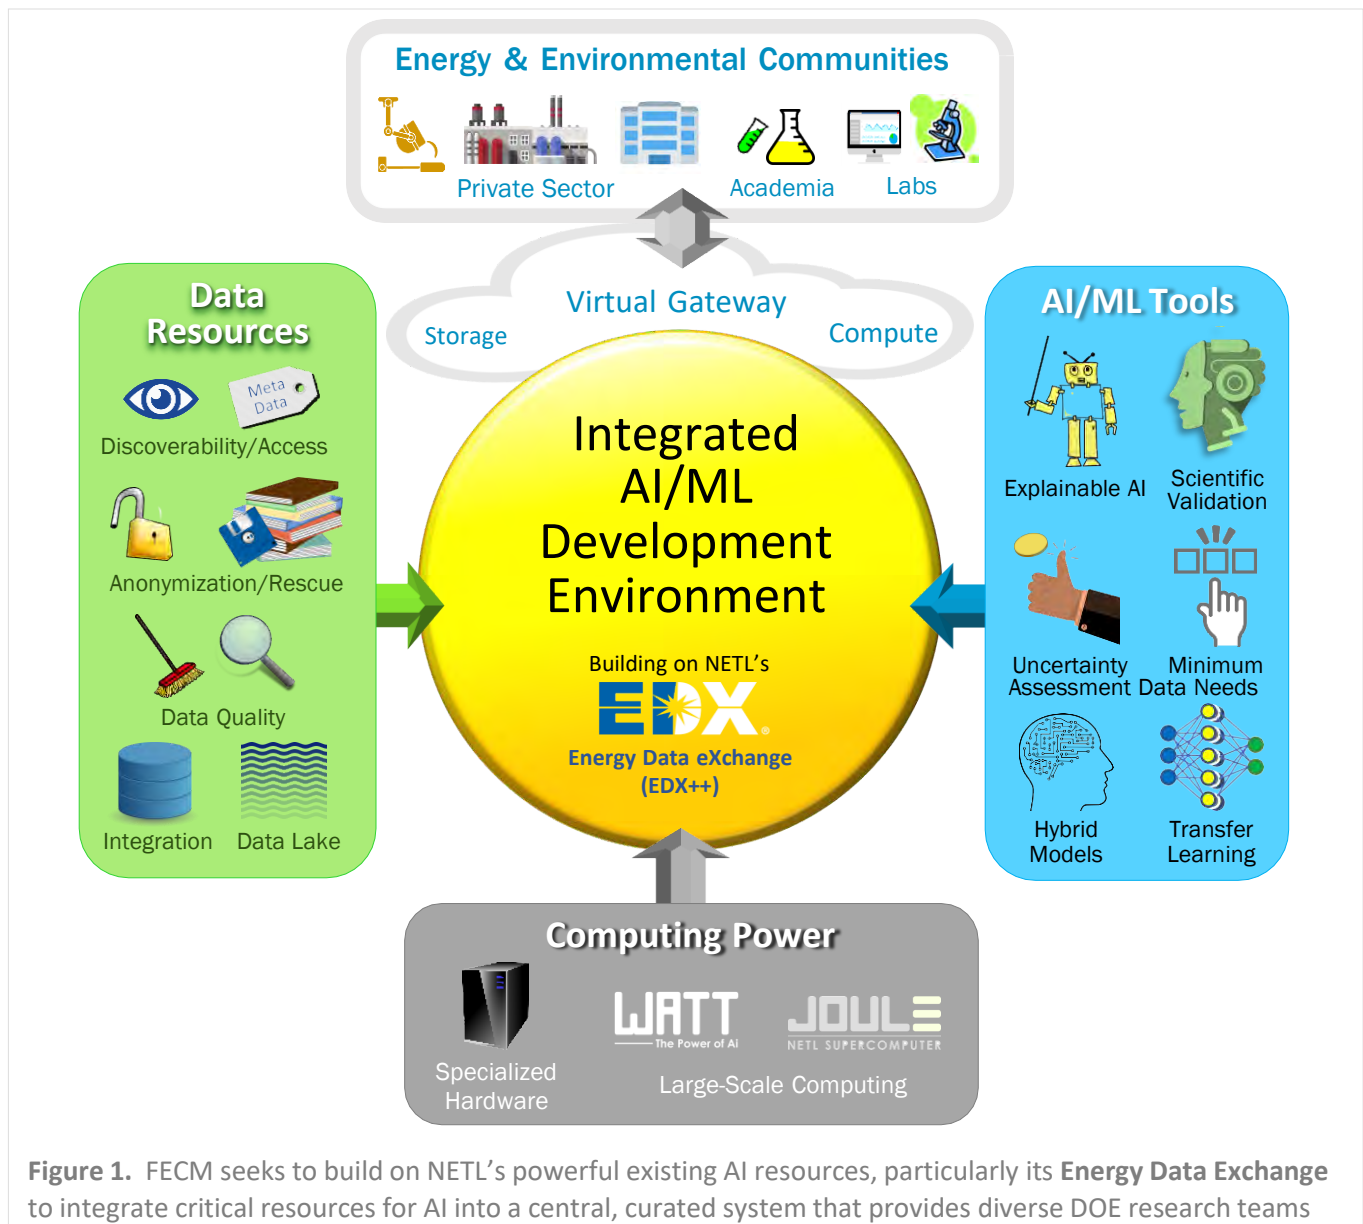


Figure 1. FECM seeks to build on NETL’s powerful existing AI resources, particularly its Energy Data Exchange to integrate critical resources for AI into a central, curated system that provides diverse DOE research teams

FECM’s AI Guidance Team recommends explicit R&D activities and other actions that FECM can pursue to accelerate larger AI benefits. These findings align with those of the recent Secretary of Energy Advisory Board (SEAB) [Final Report on AI/ML](#), which calls for a *DOE-wide*, integrated AI capability, among other resources. A draft final report by the National Security Commission on Artificial Intelligence²

similarly stresses the need for such a resource at the national level. The proposed integrated AI/ML resource at NETL could expand over time to help support a progressively larger (DOE-wide) AI/ML development environment.

Building an integrated AI/ML development environment will provide current and future FECM researchers ready access to vast, curated data sets; state-of-the-art AI tools; specialized computational resources; and a safe, collaborative environment (fully compliant with all applicable cybersecurity protocols and practices) in which to develop, test, and optimize products or solutions. Recommended R&D to create this resource is outlined below in each of the three categories shown in Figure 1.

The FECM team’s R&D recommendations on AI are closely linked and often interdependent. These dependencies indicate that implementing only a narrow subset of these recommendations, rather than the entire suite, could significantly degrade overall benefits.

Data Resources

FECM/NETL has amassed extensive (billions of attributes and features, PB of datasets), unique data sets in diverse formats and storage media (including journals, mainframe tapes, and floppy discs)—and continues to generate massive amounts of new data every day. Industry, university partners, and others also hold vast quantities of raw and sometimes proprietary data that FECM could use AI to transform into a useful knowledgebase for its stakeholders.

Data analysts now often spend up to 80% of their time just locating, obtaining access to, evaluating, formatting, and/or otherwise preparing the right data for each project.³ FECM’s AI Guidance Team identified the following R&D activities to potentially flip this paradigm—so that analysts and researchers can readily identify and prepare all relevant data—accelerating research progress and ultimately yielding better outcomes.

Recommended Actions: Data Resources

Caveat: In carrying out the recommended actions below, it is imperative to actively protect and preserve all existing workflows and data sharing agreements with industry partners. Access to these essential, high-value data sets relies heavily on our partners’ trust in DOE confidentiality—a trust that must be upheld.

Challenges: Although FECM holds an enormous volume of project data that could now drive transformative AI/ML solutions, FECM researchers today find it hard to know what data exists, where it is, what format it is in, who owns it, and other critical characteristics. Robust data curation and management tools could facilitate AI breakthroughs in plant operations and maintenance; carbon capture and storage; resource prediction, extraction, and protection; and energy storage.

“One of the major scientific challenges of our time is being able to access and effectively analyze mounting quantities of data.”

Dr. Chris Fall, Director, [DOE Office of Science](#)

Discoverability/Access — All new and existing data sets should include accurate meta data to summarize basic information about the data (e.g., type, dates, ownership, format, etc.). Meta data is a feature of the widely embraced **FAIR** principles to improve the **Findability, Accessibility, Interoperability, and Reuse** of digital assets.⁴ To the extent practical, FECM can increase data discoverability by fostering broad adoption of FAIR principles and best practices across FECM-funded research. An early step in this direction is to develop boilerplate language for future **Funding Opportunity Announcements (FOAs)**, specifying partner adherence to FAIR principles. Care should be taken to avoid unduly burdening primary investigators and to protect proprietary data if the owners are not satisfied with anonymization techniques (as described below).

Once FECM has marshalled its unique and valuable data sets, the resulting searchable, curated collection will significantly improve data discoverability. This resource can raise the participation level of potential data contributors, demonstrate tool proficiency in preparing data for specific applications, and streamline storage requirements.

Data Anonymization and Rescue — AI tools tend to improve with more data, so the ability to share data across widening scales (materials, devices, plant, company, industry) should produce more robust models that can make better predictions and decisions. Steps to increase data sharing include data anonymization tools and the rescue of legacy data. At present, many utilities, power plants, industries, and other businesses may resist sharing their raw data, fearing risks to security and customer privacy.

Data anonymization tools can prevent data from being traced to its source, but the tools currently available offer a range of security levels, and the best require significant time and investment. FECM R&D could improve the security and cost effectiveness of anonymization tools to make owners of proprietary data more willing to share their data and help generate AI tools that optimize operations across industry. Similarly, FECM needs to mount a major effort to find and develop **meta data for all existing and legacy data sets**—including those stored on mainframe tapes, paper files, floppies, and other hard-to-search formats. DOE invested in the collection and retention of this data, and it could potentially help drive future AI solutions. FECM must also plan now to avoid *future* legacy issues.

Data Quality — Researchers need to thoroughly understand the quality of their data—including its lineage, robustness, and uncertainty. These properties

“Although DOE has access to vast amounts of data and can potentially collaborate with industry to gather additional data, the quality and format of data and protecting the data are issues that need to be overcome.”

[SEAB AI/ML Final Report , p. 56](#)

vary widely by data set, and quality requirements vary by application. FECM needs a range of methods and tools to determine data quality, assess accuracy, and detect flaws. Tools are also needed to rank, cleanse, reconcile, and upgrade the data. AI-based tools could potentially make this process fast and inexpensive.

⁴ GO FAIR, FAIR Principles: www.go-fair.org/fair-principles/ See also [DOE awards \\$8.5M for FAIR data](#) to advance for AI, August 10, 2020.

Data Integration — FECM needs reliable tools and protocols for managing and integrating data assets in diverse formats and from diverse sources (other federal agencies; state, local, and tribal governments; commercial enterprises, etc.). Broad use of consistent standards would significantly expand options to create rich, application-tailored data sets. However, multiple standards exist within sectors, and convergence on a common standard will take time to emerge. FECM is well-positioned to conduct research and development efforts to create software-based tools with the following functionality:

- Parse through data for consistency
- Combine multi-source, multi-format data
- Convert data to compatible formats
- Extract or compress massive scientific data sets
- Automatically upgrade/update data from the source
- Create archives to prevent loss of investment in data
- Validate models and simulation tools

Data Lake — FECM should work with its partners and data sources on the best parameters and protocols to set up a secure “data lake” containing indexes to a diverse range of data sets. FECM needs a secure data lake to facilitate raw data discoverability and access with the ability to query data owners. The unstructured data potentially provides users great flexibility and value.

AI/ML Tools

Researchers who use traditional models based on mathematics and physics typically have little difficulty in following the logic of a model. In contrast, AI models and algorithms derived from patterns identified within large volumes of data are not amenable to human interpretation. The task is intensified in deep neural networks, which autonomously learn domain features by imprinting patterns on multiple interconnected layers of simulated artificial neurons (nodes).

While AI has a proven ability to turn massive amounts of data into useful insights, the technology continues to be a “black box,” which raises legitimate concerns: Is the model biased because it was trained on inaccurate or skewed data? Is it making decisions based on faulty patterns or relationships (e.g., traceable to a sensor error)? Users need to feel completely confident in AI models to harness their potential for problem solving and system optimization.

Challenges: AI models are inherently complex, driving the need to understand and validate model outcomes. Tools and protocols are needed to identify or mitigate bias, eliminate faulty logic, and test accuracy, often with the use of improved data. Such tools are essential to lower the risk and boost acceptance of AI.

Finding the most relevant AI models for a specific project can also require an inordinate amount of time, particularly if models are isolated in diverse silos within programs or application areas. Researchers need the ability to quickly locate a model and discern its functions, optimal uses, and data requirements.

Recommended Actions: AI/ML Tools

Explainable AI — FECM needs reliable tools to explain what its AI algorithms and models are doing. The goal of ‘explainable AI (XAI)’ is to explicitly show a model’s strengths and weaknesses, steps in analysis, and the factors that most influence the outcome (see Figure 2). Emerging XAI models often incorporate an interface that offers data visualization (to explain relationships among specific data features) and scenario analysis (to

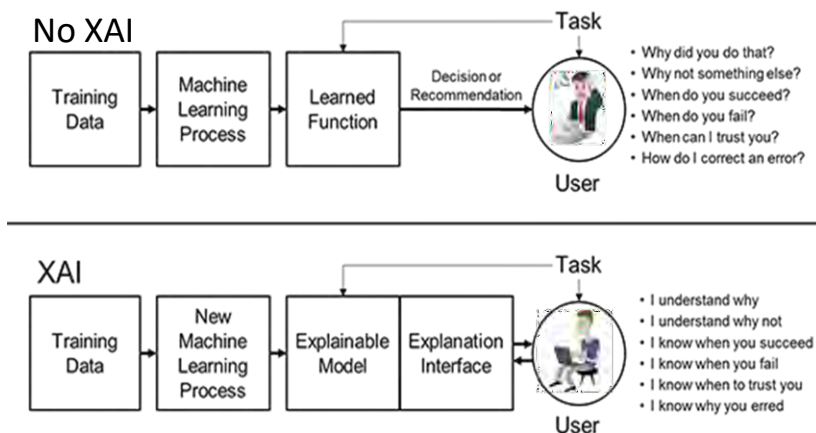


Figure 2. Explainable AI helps clarify how an AI model operates. DARPA

show how various input values affect output). Key advantages of XAI are improved debugging, faster adoption, and improved ability to audit ethics adherence.⁵

Scientific Validation — Given the science-based nature of FECM R&D, the new patterns or relationships exposed by AI must be validated by subject matter experts in the relevant disciplines. Domain experts must collaborate with AI specialists to help make sure that each model follows the science and does not violate any physical laws (e.g., the second law of thermodynamics).

Uncertainty Assessment — Users of AI models need to understand the amount of uncertainty in each model’s conclusions or other outputs. Tools that quantify model uncertainty can guide efforts to improve the model and inform the appropriate level of trust in its predictions. FECM needs to invest in an integrated tool to help industry and other partners quantify the uncertainty of a model or algorithm for a given application (similar to [CCSI](#)).

Minimum Data Needs — FECM R&D has always collected and used massive amounts of data in support of scientific accuracy. In the subsurface, for example, the tenet that the confluence of disparate geographic and geologic factors make each site unique has stimulated intensive data collection, processing, and storage; however, all of the collected data may not be significant or even useful. Research to identify the truly critical data elements could potentially save vast amounts of time, money, and effort. To the extent research can identify the critical data elements, tools would also be needed to select and more efficiently store only the useful pieces of data.

⁵ AI Multiple, Explainable AI (XAI) in 2021: Guide to enterprise-ready AI, January 1, 2021. <https://research.aimultiple.com/xai/>

Hybrid Models — As AI evolves, opportunities exist to leverage and expand model capabilities. FECM can explore the benefits of combining and augmenting models. For example, pre-trained neural nets might be combined with science-based models; to achieve this, a neural net might be trained using empirical data as well as data from a physics-based computational model.

“Those charged with utilizing AI need an informed understanding of risks, opportunities, and tradeoffs. They need awareness of the possibilities and limitations in a system’s expected performance.”

National Security Commission on Artificial Intelligence, [Final Report](#)

Transfer Learning — The initial training of a neural network typically requires intensive use of data, time, and computing power. Once trained, neural networks often require minimal additional training to become useful in applications that are similar to—yet distinct from—the one for which they were originally created. This ability to transfer the learning from an existing neural network can significantly expedite AI results and lower costs for new applications. Curated tools may be developed to help identify previously trained neural networks that are potentially relevant to an application.

Computing Power

Even as advances in AI yield impressive results, concern is growing over the availability, power consumption, accessibility, and cost of today’s supercomputers, which are used to train deep neural networks. The rapid growth in AI computing requirements raises questions about access to high performance computing (HPC) capacity. Efforts to expand AI must consider equity and other impacts. FECM maintains top-tier computer resources at NETL and recently created the *HPC for Energy Innovation Initiative* to give industry and other partners the computer access they need to explore new energy technologies. Additional solutions involve the use of cloud computing, scaled models for edge computing, more energy-efficient algorithms, tools for tracking the carbon footprints of algorithms, and more efficient chip architectures. An integrated AI development environment must provide seamless access to adequate and appropriate computing power.

Challenges: AI-associated computing creates the need for improved access to high-performance hardware, greater speed, and optimized computer architectures. Associated concerns include cost and power usage.

Recommended Actions: Computing Power

Large-Scale Computing — FECM enjoys access to two (2) large-scale computers through NETL: **JOULE** and **WATT**. The Joule 2.0 supercomputer lets researchers simulate challenging phenomena and run high-fidelity modeling tools at various scales (molecules to entire plants or carbon storage formations). A talented mathematician working 40-hour weeks for 50 weeks per year would take about 55.9 billion years to do what Joule 2.0 can do in one second.⁶ The WATT computer provides large data storage and analysis capabilities that use cutting-edge algorithms

⁶ NETL, Using Artificial Intelligence in Fossil Energy R&D, April 9, 2020: www.energy.gov/fe/articles/using-artificial-intelligence-fossil-energy-rd

and harness the power of AI and ML to address previously unanswerable problems. To maintain the value of these computing resources and keep pace with evolving requirements, FECM/NETL must continuously refresh and update these systems, integrating the latest architectures and increasing capacity.

Specialized Hardware — FECM can immediately and dramatically increase the speed and efficiency of its neural net training by acquiring new, specialized hardware. As an example, [Cerebras'](#) powerful and compact new hardware is designed specifically to train complex neural networks. Cerebras' giant chip performs this task considerably faster than large, general-purpose computers like JOULE. It also takes less space and uses far less electricity. The unique architecture of this AI hardware will provide NETL the opportunity to develop specialized software, which translates into the following:

- Faster and cheaper development of AI/ML models for real-time control and optimization of carbon capture systems, industrial plants, and power plants, improving fuel flexibility, reliability, and emissions reduction.
- Rapid development of physics-based models used for computational fluid dynamics or for the optimization or troubleshooting of power plant devices and industrial systems.

Integrated AI/ML Development Environment

Efficiently implementing the R&D activities described above and incorporating the results into NETL's existing Energy Data Exchange (EDX) network will give FECM the expanded development environment it needs to vastly accelerate AI/ML solutions *across its mission space*. As envisioned, this integrated environment will provide authorized AI researchers rapid data discoverability and streamlined access to robust data sets; tools and models; custom, secure workspaces or sandboxes; computing power; and cloud-based resources—all in compliance with federal cybersecurity requirements.

NETL is a natural choice to host this integrated AI/ML development environment. The Lab is an active participant in cooperative AI research activities across the National Laboratory System (NLS), provides specialized curation and access to large scientific data sets through EDX, serves as the lead DOE facility for geo-spatial data; and cooperatively shares its AI expertise, data tools, and computer resources among diverse FECM/NLS research communities. As a government-owned, government-operated (GOGO) lab, NETL is well positioned to serve as a user facility for all of FECM (with potential for future expansion), expediting progress on AI/ML-assisted research. The proposed development environment will be designed to:

- Preserve current interfaces for distinct FECM research communities and programs
- Facilitate access to the necessary data, software, and computing power
- Ensure automatic compliance with all applicable federal cybersecurity requirements
- Provide tailored sandboxes and expert IT support
- Reduce duplication of effort and streamline routine activities.

As host of the proposed AI/ML environment, NETL will continue to work with its sister labs and strengthen NLS foundations for advanced AI, physics-informed ML, deep learning, and neural

networks. The following important FECM/NETL/National Laboratory programs are a few of the many research communities that can leverage this integrated environment to accelerate progress in key areas:

- **Materials discovery** (e.g., [eXtreme MAT](#)). Improving the efficiency and economics of materials for extreme environments/energy applications, such as:
 - Clean hydrogen production
 - Direct air capture
 - Advanced energy storage
 - Critical minerals
- **System optimization** (e.g., [Institute for the Design of Advanced Energy Systems](#)). Supporting carbon-neutral systems in power plants and industry.
- **Offshore risk reduction** (e.g., [Offshore Risk Modeling](#) and related tools). Providing [award winning](#) data, tools, and science-based techniques to evaluate risks and security gaps in offshore hydrocarbon systems, including metocean and geohazards, analyze aging infrastructure to support the nation’s offshore spill prevention environmental and social justice goals.
- **Subsurface prediction** (e.g., [National Risk Assessment Partnership](#) and [Science-Informed Machine Learning To Accelerate Real-Time Decisions](#)). Reducing risks and improving the efficiency and environmental sustainability of subsurface exploration and carbon storage.
- **Carbon capture technology** (e.g., [Carbon Capture Simulation for Industry Impact](#)). Providing a virtual learning environment to accelerate the deployment of complex engineered systems in power and industrial applications



Recommended Actions: Integrated Development Environment

Integrated AI/ML Development Environment — As currently conceived, a qualified user could gain access to all the resources in FECM’s Integrated AI/ML Development Environment from their preferred development environment, such as [IDAES](#) or [MFI](#), via an application programming interface (API). Once connected, users will be able to:

- Search indexes in the data lake to find data sets of potential interest and pull them into the workspace, observing applicable owner requirements (if any).
- Search for algorithms, models, trained neural networks, or other software suitable to the application and pull them into their workspace.
- Schedule time on computing resources to run the software.

The sought-after data sets, models and tools, and computing resources may physically reside at NETL, within the NLS, on a Git site, on the cloud, or elsewhere.

As envisioned, approved users will eventually also have seamless access to anonymized raw data made available by external partners, including industry (following a formal permission process). In exchange for the data, these organizations may access continuously improved

algorithms or models suitable for in-plant or on-site use to optimize operations and efficiency (edge computing).

Virtual Gateway — NETL will develop a virtual gateway to the integrated AI/ML development environment to streamline entry while meeting all federal cybersecurity conditions and protocols. Building on the current portal to the EDX database, NETL will streamline tiered access by researchers who previously passed requisite cybersecurity checks and received approval. The resources to which they have access will be controlled automatically in accordance with clearance and project need.

Staff Expertise — NETL must have a highly qualified team to handle growing system capabilities and rapidly evolving cyber threats. Launching and maintaining a safe, efficient, and secure integrated AI/ML development environment will require close coordination among a team of experts in three areas of specialization:

- **Systems Administration.** To effectively manage the configuration, upkeep, and reliable operation of the integrated, multi-user computer system.
- **Cloud Systems.** To manage integration of increasingly complex cloud services with existing data, software, computing, security, and research requirements.
- **Research Science and Cybersecurity.** To maintain full compliance with federal cybersecurity rules and assure ongoing protection of hardware, software, and data systems from cyber threats while preserving required functionality for research.

Git Account — Setting up an enterprise-level account on a Git hosting service will give FECM researchers and other authorized users access to tools on other Git sites and allow them to catalog and store specialized models and other software on a single, easily accessible site. To store a new or improved software tool, users may apply to upload it to the Git. Once approved and posted, all users could gain access to these tools.

Summary

The recommendations outlined in this paper are designed to address gaps and supply the components needed so that FECM can rapidly harness the full power of AI to solve the pressing challenges now facing the energy sector and the nation. AI's potential to optimize systems and predict events or outcomes can expedite DOE progress toward net-zero carbon emission, social and environmental justice goals while supporting efficient innovation of next-generation energy technology and knowledge development.

The strategies and needs identified here for AI in FECM research programs are likely to complement those of other federally funded research programs. Therefore, the recommendations of this paper are provided for the RFI reviewer's consideration in support of their broader, national goals.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

NIYAM IT, Inc.

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Executive Summary

NiyamIT, Inc. (Niyam) greatly appreciates the opportunity to submit this Response to Request for Information for a National Artificial Intelligence Research Resource (NAIRR). Niyam is an **SBA Certified 8(a)** and **HUBZone** Small Disadvantaged Business, founded in 2007 by a group of consultants who shared a unique vision: a technology company steeped in an orderly process, yet driven by passion and innovation. Today, Niyam leads the way in developing and delivering mission-critical technologies using Artificial Intelligence (AI), Data Science/Analytics and Management, Geospatial Information Systems (GIS), Software Development, Agile/DevSecOps. Our solutions are proven to increase efficiency, streamline process flows, accelerate collaboration, and consistently provide breakthrough results - all while adapting and responding to the realities of shifting timelines and budgets. Niyam is ISO 27001:2013, ISO 9001:2015, and ISO 20000-1:2018 certified and CMMI-DEV Level 3 appraised. We strive to be a trusted partner to National Science Foundation (NSF) and the Office of Science and Technology Policy (OSTP), offering our commitment to delivering positive, measurable outcomes in support of your mission.

1. Task Force Considerations

A. Goals for Establishment and Sustainment

Large-scale digitization of information resulting from IT modernizations has leveraged data as a very valuable resource. Harnessing large volumes of data, processing it, and generating quick actionable insights surpasses human capabilities. Machine learning (ML) and Artificial Intelligence (AI) can extend human capabilities to model “never before” scenarios. Gaining a strategic advantage with AI will help the county create proactive solutions to large scale unprecedented complex problems like pandemics, loss of lives and property due to intensifying natural disasters, shifting geopolitical polarization, cyber and physical security threats, and volatile economies. In recognition of this fact, most government agencies have established AI advisory boards or similar inter-agency groups to study the potential of AI. Funding has been provided to research groups from multiple universities to develop prototypes.

The **primary goal** of the National Artificial Intelligence Research Resource (NAIRR) should be to mobilize the AI groups in government agencies to create a centralized, unified “think tank”. Such joint effort will be cost-effective, prevent duplication of efforts, and enable collaboration on a larger scale. The table in appendix A provides details on the probable transformative role of AI in each government agency. AI research is a data-driven effort, requiring specialized compute infrastructure, ML expertise, IT support staff, and strong governance support. NAIRR’s unified research platform will serve as a convergence of the best minds in the academia and varied data sets from member government agencies, compute infrastructure purpose-built for AI/neural/deep learning computations in a highly secure environment and governed by top policymakers. **Long term sustainment** of NAIRR largely depends on the collaborative efforts of government agencies and their collective drive to succeed.

Quantifying AI research success with metrics is challenging and nuanced, except in cases of genuine breakthroughs. Some platform-centric metrics of NAIRR can be indicative of success:

- No. of individuals, institutional memberships, and workspaces on NAIRR, indicating user engagement and growth.
- No. of deployable AI modules produced, or no. of milestones, indicating the progress of research objectives.
- Feedback scores from grantors for research projects indicating fair and appropriate use of the platform.

Niyam's AI/ML Solution
<ul style="list-style-type: none"> • Niyam developed custom-built AI/ML tool Flood Assessment Structure Tool (FAST) to compute flood risk, increasing speed and accuracy. • FAST analyzes NYC 800,000 structures for 100-year flood losses in less than 8 seconds (previously it was >10 hours), 275x faster than existing traditional software.

B. Ownership and Administration

An AI working group established by the NAIRR Task Force should oversee the operations of AI subgroups, established in each participating government agency. These agency AI subgroups should be responsible for research initiatives piloted by the parent government agency. The AI subgroup can be owned by the participating agency. This agency subgroup will define research problems and invite new ideas/proposals for transformative AI to be published as grant opportunities that will be made available to accredited universities. After successful evaluation, the grants administration should be handed over to an agency similar to grants.gov.




Figure 1. Ownership and Administration of NAIRR

Responsibilities of agency-specific AI subgroup:

- Publish research opportunities on the NAIRR platform for aspiring research groups
- Evaluate the feasibility of submitted proposals and grant approvals
- Disburse funds for NAIRR to the AI working group with expenses calculated based on usage requirement of NAIRR platform that includes data sets, algorithms, and compute times. (This replaces the need for government agencies to fund individual research groups)
- Track research progress report to AI working group

- Identify platform insufficiencies such as compute requirements and unavailability of data, connectivity, availability, and suggest improvements to the NAIRR working group.

C. Model for Governance and Oversight

Governance and oversight objectives of the AI working group should be to enable the use of high-quality data for ethical AI research on a secure and robust platform while protecting privacy and adhering to the compliances as applicable in each business domain. Governance should not be a blocker, but an enabler to innovation. A balance between governance and enablement must be established to foster innovation. Research projects on NAIRR should be given autonomy for exploration and experimentation. In turn, they need to provide 100% visibility into the data sets being utilized, the purpose of the research, and adherence to privacy and compliance laws. The 5 anchors for oversight and governance are:

- **High-quality data:** The U.S Open Government Directive required that all agencies post at least 3 high-value data sets online and register them on data.gov within 45 days. As an extension of this directive, the NAIRR working group should procure fresh data sets from participating government agencies regularly.
- **Ethical AI:** Research projects on NAIRR would have the potential to have a high impact on the daily lives of the public. Ongoing scrutiny is crucial to ensure that the NAIRR is committed to enabling unbiased and trustworthy AI research.
- **Compliance & Privacy:** A unified charter of privacy and compliances requirements (e.g., HIPAA, SOC, PCI) should be jointly maintained by member government agencies. All AI research projects on NAIRR should fulfill these requirements for continued access.
- **Security:** Security threats to NAIRR exist at dual levels, at data repository level and programmatic/model level. Many AI systems could be composed of open-source libraries with security vulnerabilities. There exists a possibility of “inversion attacks” where an AI model can be hacked, revealing information about itself and the data that it was trained on. Such threats should be identified and monitored to prevent misuse of the NAIRR.
- **Platform Infrastructure:** Factors that lead to high-performing AI platforms are high computing capacity, storage capacity, and networking infrastructure. NAIRR working group should monitor and adapt the platform infrastructure to meet the needs of the research community.

D. Creating and Maintaining Shared Computing Infrastructure

Resources required for data storage/preparation are different in scale and nature from those required for AI computation. Therefore, infrastructure for storage and computation can be decoupled from each other.

Data storage and processing requirements: The NAIRR platform will have to plan for a massive storage repository of data needed for ML model training, and a high velocity of incoming data streams for model inference and predictive analytics. Data sets for ML and AI can reach hundreds of terabytes to petabytes. Data consolidation from multiple

sources, selection, and preprocessing, such as filtering, categorization, and feature extraction which are the primary factors contributing to a model's accuracy and predictive value are significant factors that will influence storage and processing infrastructure.

Computing resources: The technology that powers AI is unique and always evolving. Specialized hardware, such as GPUs is critical for machine learning (ML) and subsequently AI. NAIRR should evaluate the pros and cons of a “build vs buy” approach. Enterprise AI has multiple success stories that feature high-level ML and Deep learning (DL) services like AWS SageMaker, Azure Cognitive Services, Google Cloud Machine Learning Engine, H2O.ai, IBM Watson Studio, and ML, etc. Among the multiple choices available for architecture and deployment for AI, hyper-converged infrastructure (HCI) systems offer the most density, scalability, and flexibility. Most AI systems run on Linux VMs or as Docker containers. Some popular AI frameworks and many sample applications are available as prepackaged container images from Nvidia and other vendors. Some of these applications include Computer vision such as image classification, object detection, image segmentation, and image restoration, speech and natural language processing, recommendation systems that provide ratings or products based on prior user activity, content analysis, filtering and moderation, pattern recognition, and anomaly detection. These applications can be used in a variety of business domains like fraud analysis, surveillance systems for physical security, geologic analysis for oil, gas, and other forms of energy, cybersecurity, and automation, etc.

The system components most critical to AI performance are:

- **CPU** - Responsible for operating the VM or container subsystem, dispatching code to GPUs, and handling I/O. E.g. Second-generation Xeon scalable platinum or gold processor, AMD Epyc CPUs.
- **GPU** - Handles ML or DL training and inferencing (ability to automatically categorize data based on learning) E.g., Nvidia P100(Pascal), V100(Volta), or A100(Ampere) for GPU training and V100, A100, or T4 (Turing) for inference.
- **Memory** - AI operations run from embedded high bandwidth GPU memory that is much faster than conventional DRAM.
- **Network** - AI systems are often clustered together to scale performance and are connected with 10Gbps or higher Ethernet interfaces. They can also include InfiniBand or dedicated GPU (NVLink) interfaces for intra-cluster communications.

Portal-based educational Tools and services: NAIRR can provide members access to prebuilt algorithms. The portal can also support virtual events to host prototype presentation opportunities regularly. Prototypes can be scored based on technical merit and award grant opportunities to winning scores. “Citizen scientist programs” can be initiated to generate interest among aspiring entrepreneurs. Educational seminars and certification tracks can be developed as revenue-generating resources. Video tutorials, wikis, blogs, and chatbot-based help and forums powered by an expert knowledge base will help researchers achieve their goals.

E. Barriers to High-quality Government Data Sets

Government agencies are using data dissemination and related communication strategies to extend the utility of their data to a wide audience. The effectiveness of these

strategies depends on the characteristics of data, target audiences, channel of dissemination, and data formats. A successful approach would need to consider the different needs and sophistication of data requirements of the research communities. This approach should also define the data sharing formats and channels for impactful use in AI research.

Barriers to dissemination are:

Knowledge barriers- A common barrier to the use of publicly held data sets and sharing of privately-held data is a lack of knowledge and awareness about the data available and methods of accessing it. It can be overcome by publishing the availability of data sets and accompanying lineage.

Technical obstacles- Technical obstacles are due to incompatibilities in data formats and preparation methods. To overcome these, data must be transformed, organized, and restructured into a commonly acceptable format using a common methodology. Though technically feasible, the process is elaborate and time-consuming.

Documentation barriers – Differences in data documentation, such as undocumented codes, coding conventions, or missing documentation can result in data that can only be used with difficulty or is completely unusable. This can be overcome by adopting standards for cataloging and documentation.

Conflicting values and obligations – Intellectual property and confidentiality concerns are often cited as reasons for not sharing data. Also, premature release of data might allow some other organization to publish first, and any sharing could deprive the original data collector of longer-term opportunities to mine the data. Enforcing data sharing standards across all NAIRR members will ease data access.

F. Security Requirements and Access Controls

Researchers of accredited and approved universities, which are members of NAIRR should be granted access to the NAIRR platform on a project-by-project basis. They should be required to undertake training in privacy and ethics to become eligible for access, receive certification and meet security requirements. To reduce security risks, the AI subgroups should formally review the research outline and build data sets for analysis using minimum no. of variables and provision minimum compute resources required for the project. Any articles, papers, or materials developed during the research should be examined and cleared by the subgroup prior to publication. The platform should easily allow for the discovery of users, roles, permissions, and access logs.

G. Privacy and Civil Rights and Civil Liberties

Anonymization of data sets is the first measure to ensure the privacy of data subjects. There could also be several other indirect identifiers that could lead to deductive disclosure such as small sample sizes or unusual characteristics of occurrences. Samples taken from specific sub-populations, geographic areas, and linked data sets can be a challenge when trying to protect subject identities. NAIRR can consider data suppression strategies to minimize privacy risks. As data and AI research findings take a variety of formats during the project from system to system, maintaining an audit trail will increase privacy confidence in the NAIRR platform.

H. Sustainment of NAIRR

Partnerships with the corporates can be considered as a source of sustainment for NAIRR. Access to the NAIRR platform can be configured as two types, for government agencies and the commercial sector. Government agencies or member subgroups will fund NAIRR from a portion of research grants. The commercial sector will fund NAIRR for using the platform for data-as-a-service and compute-as-a-service.

I. Parameters for Establishment and Sustainment of NAIRR

Under the AI working group, we propose agency roles for 5 anchors of governance:

Chief Data Officer (CDO): To ensure data quality and promote data governance. Duties include:

- Maintain high availability and quality of data.
- Oversee data governance.
- Create and maintain a data management system to secure data collection and preprocessing of data to promote usability.
- Foster data sharing across the government and industry and establish open data initiatives.
- Identify potential AI research areas and develop implementation plans to acquire or collect data for future research

Chief Ethics Officer (EO): To promote responsible and ethical use of AI on the NAIRR platform. The duties of the EO include:

- Ensure ethical procedures are implemented and adhered to by NAIRR users.
- Examine ongoing research on the platform to enable unbiased and trustworthy AI research.
- Raise ethical considerations for proposed research areas.
- Install an ethical and responsible AI culture interagency as well as across the government and industry.

Chief Privacy & Compliance Officer (CPO): To uphold compliance and privacy requirements. The duties of the CPO include:

- Assess and manage risk related to privacy considerations and compliance.
- Examine ongoing research and future research for information privacy.
- Oversee and monitor implementation of a compliance program within NAIRR.
- Promote adherence to privacy laws within NAIRR.

Chief Information Security Officer (CISO): To protect NAIRR users, assets, and IT systems from security threats. The duties of the CISO include:

- Prevention, investigation, and handling of security threats.
- Establish security practices as well as implement security systems.
- Uphold adherence to governance related to security.
- Responsible for disaster recovery and continuity of operations.

- Assess and handle the risk of cyber threats, data loss, and fraud.

Chief Infrastructure Architect (IA): To maintain computing resources and IT infrastructure. The duties of the IA include:

- Design and implementation of NAIRR enterprise infrastructure.
- Promote and maintain platform infrastructure quality and high availability.
- Improve customer experience of NAIRR IT systems.
- Explore and lead customization and modernization efforts.

2. Prioritizing NAIRR Capabilities and Services

AI being a data-driven initiative, NAIRR should prioritize the creation, collection, preservation, storage, retrieval, and distribution of machine-readable data that can fuel a wide spectrum of AI research.

Creating a unified data platform as a first step will give an idea of the variety and volume of data workloads, which will help form a strong foundation for architecting a robust compute platform for AI and ML processes. Domain experts should be onboarded in this stage for

data sourcing, management and cataloging curated data sets. NAIRR should publish the cataloged metadata for every data set. Regular refresh and archival cycles on existing data sets (and metadata) should be scheduled, to maintain relevance.

Niyam's Unified Data for FEMA

- We managed over 480 TB of structured and unstructured data from various sources for combined business intelligence for Risk Management Directorate (RMD).
- We developed ETL workflows for data ingestion and advanced analytics capabilities utilizing modern data science techniques such as clustering, Bayesian, and other statistical models.

3. Reinforcing Ethics and Responsible Research

NAIRR and its components can reinforce principles of ethical and responsible research and development of AI by enforcing the use of the Responsible AI framework for all projects implemented on the platform. Within this framework, all ML models should be comprehensive, explainable, ethical, and efficient.

- **Comprehensiveness:** The AI model has clearly defined testing and governance criteria
- **Explainability:** The purpose, rationale, and decision-making process of the AI model can be understood by the average end-user
- **Ethical:** The AI initiative has processes in place to seek out and eliminate bias in ML models
- **Efficient:** The AI model can run continually and respond quickly to changes in the operational environment.

At the pre-design stage, the agency subgroups should evaluate research requests on NAIRR by a scrutiny of the problem documentation that includes:

- a. Business context of the research problem undertaken
- b. Business justification for the algorithm to be developed

- c. Model parameters are used for tuning the model to maximize its performance without overfitting or creating a high variance
- d. Feature choices (inputs) and definitions (outputs)
- e. Any customizations to the algorithm if it was reused
- f. Instructions for reproducing the model
- g. Examples for training the algorithms and datasets used
- h. Examples for making predictions from the algorithm

After successful evaluation, the research projects should be developed within the guidelines of Responsible AI as below:

Shared code repositories: Shared code repositories facilitate efficiency by eliminating rework and reducing the processing overheads of the compute platform. Researchers can reuse existing models/algorithms as stepping stones to further their research on solving newer problems.

Approved model architectures: New model architectures should be approved by the NAIRR working group by evaluating them on explainability and interpretability. This is an important factor to eliminate issues related to fairness, bias, transparency, and accountability.

Sanctioned variables: Datasets made available of research should not contain any personally identifiable information (PII) directly or indirectly. Each dataset should be tagged with its summary statistics indicating the distribution of values, to eliminate bias.

Established bias testing methodologies to uphold fairness, civil rights, gender equity in the models created for AI systems.

Stability standards for active machine learning models to make sure AI programming works as intended and does not cause memory leaks and performance bottlenecks for the platform.

Implementing Responsible AI: The most important catalyst for solid governance for implementing Responsible AI is model validation and reproducibility. Model validation is the process of ensuring that the AI model is performant, statistically sound, delivers statistically significant benefits, and meets the definition of “success” put forward by the AI project.

Researchers should group their model by the project. Each attempt to train a model for that project is called a “run,” with all the runs for that project being rolled up into an “experiment.” Putting forth a simple metadata framework centered on the concept of an experiment yields increased visibility and auditability for any AI project. Metadata necessary to reproduce an experiment or a run of an experiment:

- a. Type of algorithm used for the development of the model
- b. Features and transformations used in the model
- c. Data snapshot or identify the data set used
- d. Model tuning parameters
- e. Model performance metrics
- f. Verifiable code location from source control management

g. Training environment setup used for model training

To test the validity of the models, the model should be tested on these behaviors:

- It achieves acceptable statistical performance for a sensible offline metric (accuracy)
- It achieves a statistically significant improvement when compared to control on some online metric or key performance indicator (KPI) (Clicks, conversions, purchases)
- It is statistically sound, there is no data leakage, and the supervised ML problem was framed correctly.
- The performance of the model can be successfully explained based on available features.

4. Existing AI Building Blocks for NAIRR

As the potential of AI research is being recognized, many government organizations are developing programs, resources, and services. NAIRR can establish partnerships to use these as building blocks.

Open data initiatives that provide data sets for download:

- Data.gov – federal government data repository for publicly available federal, state, local, and tribal government information.
- Census.gov – provides United States Census data publicly.
- Data.gov.uk – United Kingdom’s published data by their government.
- Open Knowledge Foundation Global Open Data Initiative – provides a repository of the world’s open government data publications.
- Federal Reserve FRED Economic Data – provides economic data for research.

Compute and data analytics management resources and AI platforms:

- National Science Foundation’s (NSF) Advanced Computing Systems and Services (ACSS) program.
- IBM Watson Studio, RapidMiner, Alteryx, MATLAB, Tableau Server, RStudio, Qlik Sense, Google Cloud AI Platform, Azure Machine Learning Studio.
- Veritone aiWARE Government platform.

Organizations promoting AI through advocacy and innovation efforts:

- Department of Energy (DOE) Artificial Intelligence and Technology Office (AITO).
 - First Five Consortium – a collaboration with industry to promote AI capability, Niyam is a partner of the First Five Consortium.
- United States Department of Agriculture (USDA) AI Institute for Next Generation Food Systems (AIFS) – using AI to drive food systems in the USA.
- USDA National Institute of Food and Agriculture (USDA-NIFA) and NSF.
- Department of Commerce National Artificial Intelligence Advisory Committee (NAIAC).

5. Public-Private Partnerships for NAIRR and Exemplars

NAIRR should provide a separate public and private tier of services, for sustainability. For the public platform, partnerships should be developed with prominent research institutions and industry. Examples of such collaboration include:

- a. Partnership between Niyam and Pacific Disaster Center (PDC). We collaborate with the University of Hawaii and other research institutions to research mitigation steps for risks during natural disasters
- b. Federal Emergency Management Agency (FEMA) provides publicly available flood products to reduce flood risk under the Risk Mapping, Assessment, and Planning (Risk MAP) program.
- c. DOE has partnered with AI-capable companies in the industry, including Niyam, to focus on Humanitarian Assistance and Disaster Response.

The private tier of services of NAIRR can host opportunities like coding competitions to raise funding for sustainment. For example, Kaggle has a business model based on partnerships with private companies to host competitions.

6. Limitations in Democratize Access to AI R&D

Observability, Visibility, and Control: Democratizing AI across a broad spectrum of government agencies, researchers, and citizen scientists will lead to heavy usage and can cause performance degradation, network congestion, data store deadlocks and contention, and other resource allocation complexities. These issues can be addressed by a holistic continuous performance monitoring of the platform and individual models served in real-time. This requires first-class integration with dashboards and visualization software, that can generate usage reports, alerts, and risk mitigation steps.

Intellectual Property rights: Researchers using the NAIRR platform will be on a path to develop powerful AI tools, with novel ideas. The perceived benefits of democratization may not be achieved without decisions about who owns the intellectual property rights. NAIRR's working group should have a strong legal framework that addresses rights and responsibilities around patents, copyrights, and trade secrets.

Cost visibility and management: The ownership and administration structure of NAIRR as explained in Question 1-B, detailed financial reporting for federal stakeholders will be required. Costs associated with the development and maintenance of NAIRR are likely to fluctuate based on a no. of different factors. Therefore, building visibility of costs using responsible accounting strategies is important. Chargeback and show-back strategies can be used to shift responsibility to participating agency subgroups or members and encourage them to become more aware of costs. They will also be helpful in budgeting, planning, and forecasting.

Compliance attestation (SOC 2, HIPAA): Any compliance failure could put the NAIRR platform at massive PR and financial risks. The agency subgroup should ensure that all AI research within their jurisdiction is compliance attestable and there are no violations.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Representatives from the National Oceanic and Atmospheric Administration (NOAA) Artificial Intelligence Executive Committee (NAIEC) and the Center for Artificial Intelligence (NCAI)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

RFI Response To: OSTP and NSF “Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource” issued 23 July 2021 [1]

Date: Thursday, 30 September 2021
From¹: Representatives from the National Oceanic and Atmospheric Administration (NOAA) Artificial Intelligence Executive Committee (NAIEC) and the Center for Artificial Intelligence (NCAI)

Contact: Dr. Rob Redmon, NOAA/NESDIS NCAI Lead

Background

The contributors to this RFI response are representatives from the NOAA Artificial Intelligence Executive Committee (NAIEC) and the NOAA Center for Artificial Intelligence (NCAI) that is authorized via the National AI Initiative Act of 2020 (DIVISION E, SEC. 5001) [2]. Through becoming law on January 1, 2021, the initiative provides a coordinated authorization for the entire Federal government to accelerate Artificial Intelligence (AI) research and application for the Nation’s economic prosperity and national security.

NOAA has developed several strategies to harness the potential of emerging technologies, including the NOAA AI strategy [3] and is unlocking the full potential of Cloud, Data, and AI/Machine Learning (ML) capabilities in support of its mission.

Response

{Italicized Times New Roman text in this section is taken verbatim from the RFI notice.}

1. "What options should the Task Force consider for any of roadmap elements A through I, and why?"

Response: All elements listed are important to the development of a high quality roadmap. We have provided suggestions for consideration to items:

- *"A. Goals and metrics for success,*
- *D. Access to curated data sets and educational tools,*
- *E. Dissemination and use of high-quality government data sets,*
- *G. Civil rights and civil liberties requirements,*
- *I. Agency roles and responsibilities."*

A. "Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;"

Response: While intended for NOAA’s mission areas (e.g. ocean observations and exploration, weather, climate, etc), the goals from NOAA’s AI strategy could be mapped and leveraged broadly. Below are a few ideas for the NAIRR to consider against their use cases.

Submitted 30 September 2021

- Goal 1: Establish an efficient organizational structure and processes to advance AI
 - Conduct gap analysis of agency, academic, and industry needs and readiness to support and leverage a NAIRR (e.g., with data sets and services, educational materials, and open source development tools),
 - Leverage existing OSTP to United States Government (USG) agency Working Groups to build cross agency coalitions to share data, expertise, and resources,
 - Encourage agency staff exchanges to the NAIRR to capture broad perspectives including from NOAA,
 - Success metric: Breadth of USG, Academia, and Industry engagement
- Goal 2: Advance AI research and innovation, and research to applications
 - Leverage existing agency Federal funding mechanisms to develop competitive grants, X-prizes, Cooperative Research and Development Agreements (CRADAs), Hackathons, Other Transactional Authorities (OTAs), etc.,
 - Collaborate with Industry “AI for Good” activities,
 - Create AI-ready data, Trustworthy, Explainable, Equitable, and computationally efficient AI/ML measures and standards to infuse as requirements for Federal funding opportunities,
 - Develop interoperable Federal data lakes—Cloud optimized repositories with structured AI-ready data and baseline open source development support—of information products with free upload/download (aka egress),
 - Success metrics:
 - (a) Diversity of stakeholders involved including opportunities for students and historically underrepresented groups;
 - (b) Diversity of novel research activities supported and their potential mappings to economically valuable and/or decision-making applications;
 - (c) Economic value across sectors from applications transitioned to an Operational, Decision Support Service or Commercial setting;
 - (d) Federal efficiencies realized from implementing improved workflows
- Goal 3: Strengthen and expand AI partnerships
 - Develop USG agency level partnerships
 - Create a USG, International, Academia, Industry Advisory Working Group,
 - Foster and enhance partnerships and synergies with enabling technologies and activities (e.g., scientific data stewardship, commercial Cloud and AI pipelines),
 - Encourage agency, academia and industry staff exchanges,
 - Develop guidance for the treatment of intellectual property rights (IPR) resulting from AI partnerships, with an aim towards ensuring open approaches to AI development,
- Goal 4: Promote AI proficiency
 - Encourage the emergence of domain and domain agnostic Communities of Practice,
 - Encourage agencies, academia, and industry to collaborate on curated libraries of “Learning Journey’s”,

- Support proficiency in AI/ML via co-productive and competitive hackathons, certificate and higher educational programs, leveraging partnerships in Goal 3.

D. "Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;"

Response: The following capabilities should be considered:

1. Extend the ai.gov web accessible portal with features including AI-enabled Application Programming Interfaces (APIs) for researchers to search for and discover curated AI-ready datasets, educational materials, workshops and other engagement opportunities,
2. Develop and manage a Cloud resource pool with AI ready development sandboxes that users can apply for time to use,
3. Develop interoperable data lakes of curated Federal datasets, with baseline open-source development tools leveraging open repositories (e.g. Git). These curated data should be AI-ready including machine actionable documentation, readily working software tools, and adhere to Findable Accessible, Interoperable, Reusable (FAIR) standards,
4. Develop platform to include interactive notebooks (e.g. Jupyter or similar) as well as an interactive forum for exchanging community ideas and resolving technical challenges (e.g. Slack or similar),
5. Develop a framework to assess the AI-readiness of government produced data sets.

G. "An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research;"

Response: The NAIRR should facilitate the development of Trustworthy AI standards and metrics to identify and minimize bias, and infuse these standards as required elements in funding opportunities. The NAIRR should include a mechanism for capturing biases discovered in operational products, which could be used to spur additional evaluation and development.

I. "Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities."

Response: OSTP should consider developing new or leveraging existing OSTP to USG agency Working Groups with representation from each agency; agencies contribute to national AI-ready data standard, and Trustworthy AI metrics.

In addition, agencies will need to devote significant effort to evaluating and improving the AI-readiness of their open data assets, ideally coordinated through agency Chief Data Officers. Establishment of a National Research Resource must acknowledge and support this crucial data

curation, and standards metrics development work.

2. "Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?"

Response: We recommend consideration of all capabilities suggested in item D above with particular emphasis on the creation of AI-ready development sandboxes connected to Federal data lakes of curated AI-ready data and baseline open-source development support.

3. "How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?"

Response: Trustworthy and Explainable AI (XA) best practices should be developed, training promoted and assessments performed to identify and minimize bias throughout the Research to Operation/Application/Commercialization pipeline. In particular the NAIRR should support the development of diverse teams assuring focus on AI issues.

4. "What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?"

Response: The NAIRR could leverage and extend its prior efforts to develop AI readiness matrices and could capitalize and support existing agency level strategies and developing implementations such as NOAA's AI, Cloud and Data strategies, and new Center for AI (noaa.gov/ai), as well as the recently launched National Science Foundation's National AI Institutes Program [4].

5. "What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?"

Response: Participating agencies could develop cross-agency pollinating activities: (a) announcements of opportunity, (b) CRADAs, (c) OTAs (e.g. [NOAA's OTA with Google](#)), (d) competitive X-prizes, (e) co-productive and competitive hackathons, (f) and training events. One exemplar is the NASA Center of Excellence for Collaborative Innovation (CoECI), who collaborates with innovators across NASA and the Federal Government to generate ideas and solve important problems by working with global communities via the NASA Tournament Lab. NOAA has several examples of direct experience in working with CoECI and others to conduct international competitive X-prizes (e.g. <https://www.fishnet.ai/>), including the recently concluded Model the Magnetic Field (MagNet) (<https://ngdc.noaa.gov/geomag/mag-net-challenge.html>) challenge which received thousands of potential model solutions that converged towards the best achievable performance and resulted in several prize awards to the international data science community.

¹Any views or opinions expressed herein, are those of the author(s) and do not necessarily

Another possible opportunity for public-private partnerships is in data preparation and dissemination (e.g. [NOAA's Big Data Program](#)). Private data brokers have enormous capability to clean, harmonize, and distribute data. Perhaps there is an opportunity for those brokers to achieve a public relations win by showing that these capabilities can also be used for good. For example, imagine a private data broker agreeing to develop a data infrastructure to bring together climate data from various Federal agencies, clean and harmonize the data, transform it into cloud-optimized formats, establish open APIs, and release it to the public for free as an AI-ready climate database, alongside archival of the most critical of these improvements in a national archive alongside the source data.

6. "Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?"

Response: One significant challenge along the path to democratized access to AI R&D resources may include data transfer costs from data provider repositories to the user's NAIRR processing environment. A mitigation step for this issue could be for AI R&D resources to include fully functional AI-ready sandboxes including AI-and Analysis-ready datasets and AI/ML processing tools & pipelines, as well as encouraging development using open source libraries and platforms.

Additionally, restrictive intellectual property rights (IPR) are a concern, and one mitigation step is to develop guidance for the treatment of intellectual property rights (IPR) resulting from AI partnerships, with an aim towards ensuring open approaches to AI development

Another significant challenge is today's lack of interoperability between Cloud provider architectures and data formats, inclusive of USG data products, and a mitigation step might be to include this need in the development of an AI-ready standard, along with encouraging the use of open source libraries and platforms.

References:

- [1] Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource, issued 23 July 2021, with due date extended to 1 October 2021.
- [2] National AI Initiative Act of 2020 (DIVISION E, SEC. 5001) became law on January 1, 2021
- [3] NOAA Artificial Intelligence Strategy (signed), Version 2020-09-17.
- [4] New NSF AI Research Institutes to push forward the frontiers of artificial intelligence, 26 August 2020.

¹Any views or opinions expressed herein, are those of the author(s) and do not necessarily

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Noblis

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

INTRODUCTION

As an independent science, technology, and strategy corporation, Noblis delivers innovation and management expertise from a position of independence and objectivity. For more than 25 years, we have worked across the defense, national security, intelligence, and civilian domains to solve difficult problems that help our government—and nation—operate more effectively and more efficiently. By assembling top talent across many disciplines, we apply the right expertise to support our clients' most critical missions with practical and actionable solutions.

Noblis is an established industry leader in solving the most pressing national security issues of today and tomorrow and routinely delivers algorithms, analytics, tools, and services while sitting side by side with our clients. As a 501(c)(3) company focused on basic and applied research, Noblis is keenly aware of the value of collaboration and works closely with partners from across the public and private sectors, as well as academia. Among its many partnerships, Noblis is a member of eight Other Transaction Authority (OTA) consortia to deliver prototypes, subject matter expertise, and requisite capabilities in the areas of aviation, medical, energy and the environment, counter weapons of mass destruction, information warfare, command, control, communications, and cyber, and systems of systems. As a consortia member, Noblis and its partners share ideas, training, concepts, and capabilities openly. This level of “crowdsourcing” generates enormous benefit to include vastly enhanced subject matter expertise, rapid technology innovation, and best of breed tools, products, and services.

Noblis' Artificial Intelligence (AI) experts leverage our state-of-the-art AI Laboratory to accomplish both basic and applied research. Our experts improve analyst productivity by designing, implementing, and evaluating full stack software solutions using Agile and DevOps practices with a backbone of modern machine learning (ML) and AI algorithms. These solutions allow analysts to better organize, explore, and understand terabytes of data. Our experts focus on developing and deploying analytic tools and capabilities within existing operational environments lowering client risk and enabling them to generate new mission insights, increase speed to delivery, improve confidence, and save financial resources fully and more rapidly. Examples of Noblis expertise in this important field include:

- Improving the recognition performance on unconstrained face imagery.
- Integrating algorithms for cargo imaging and sensor applications.
- Developing algorithms for face morphing and aspect ratio manipulation.
- Leveraging deep learning networks to generate captions for images and videos.
- Identifying combinations of algorithms to counter adversarial machine learning.
- Using natural language processing (NLP) to create and maintain taxonomies.
- Progressing ways to create, track, and present AI rationale and chains of evidence.

The final report from the National Security Commission on Artificial Intelligence (NSCAI) included the following statement, “*the rapidly improving ability of computer systems to solve problems and to perform tasks that would otherwise require human intelligence—and in some instances exceed human performance—is world altering*”.¹ Noblis concurs with the NSCAI that AI can have world altering impacts and is highly supportive of the goals of the National Artificial Intelligence Research Resource (NAIRR) Task Force. Our response to question 1 of

¹ NSCAI Final Report, pg. 7

the NAIRR Request for Information (RFI) below illustrates the importance of bringing together people, process, data, and technology to achieve NAIRR Task Force objectives.

- *People*: Pairing together multi-disciplinary teams comprised of data analysts, data scientists, technologists, academics, and domain experts is an essential element to ensuring that research will be achievable, impactful, and adoptable.
- *Process*: Effective governance, to include providing direction for priorities and funding, providing training, and ensuring excellent internal communication and external outreach are essential to generating quick-wins and delivering value-added research and capabilities.
- *Data*: Sufficiently structured and validated data made available for training algorithms, testing capabilities, and delivering unambiguous, repeatable, ethical, and trustworthy results is essential for generating new insights and improving confidence in findings.
- *Technology*: Leveraging cloud-based computing power and expanding the use of technology such as NLP and deep learning to enable automated data-curation and automated entity/taxonomy development of large corpuses of data are essential for increasing speed to deliver relevant research and findings.

1 WHAT OPTIONS SHOULD THE TASK FORCE CONSIDER FOR ANY OF ROADMAP ELEMENTS A THROUGH I AND WHY?

A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success.

The National Science Foundation (NSF) – led National Artificial Intelligence Research Institutes (herein referred to as Institutes), twenty now in total² represent an exciting opportunity to research, develop, test, and apply AI to challenging problem sets and to many new fields. The Institutes, which are located across the country and are comprised of academia, government, and industry are excellent AI building blocks. A strong foundation comprised of effective governance, available infrastructure, robust collaboration, and assured funding are necessary to focus ideas, generate value-added research, share best practices and lessons learned, and to ensure the building blocks are solid and sustainable for years to come.

Goal 1: Establish NAIRR Governance. Noblis recommends that a NAIRR Governance Board (herein referred to as the Board) be established and chaired by the Director of the Office of Science and Technology Policy (OSTP). The Board should include the NSF Director and the Director's from each of the Institutes (one per each). The NSF Director, in consultation with Director/OSTP should select the directors of the Institutes keeping in mind that a proper balance between academia, public, and private sector leadership is advisable to ensure optimal Board perspectives and decision-making. Noblis also recommends that the Directorship for each Institute rotate periodically (e.g., annually) between academia, government, and industry to maintain high levels of involvement, ownership, and transparency. The Board should meet monthly within the first year of initiation to establish and implement a strong and sustainable modus operandi. This too will help build networks, break down barriers, identify joint needs and requirements, and allow the members to share, leverage, and apply best practices, both individually and collectively. Significant Board functions include:

- Prioritizing and coordinating Institute research. – See Goal 2

² https://www.nsf.gov/news/news_summ.jsp?cntn_id=303176

- Funding and maintaining the underlying infrastructure (i.e., cloud technology and tools for all Institutes). – See Goal 3
- Providing an AI Technology Accelerator Platform (TAP) to track and monitor projects. – See Goal 4
- Implementing an acquisition model that pairs customers (i.e., government) to performers (i.e., Institutes) and enables funding to be put on contract expeditiously. – See Goal 5
- Providing data governance to make relevant data widely sharable and protected accordingly for research through data sharing agreements; Providing technology, data, research, or related policy recommendations; and Communicating NAIRR plans, successes, and requirements to the President, Congress, and the public as appropriate. – See Goal 6

Suggested performance metrics for Goal 1 include:

- Governance Board established / All members identified and in position
- Terms of Reference and Concept of Operations published
- Number of meetings held / Number of decisions made
- Number of performance metrics met (for goals 2-6), and associated results

Goal 2. Prioritize and Coordinate Institute Research. The range of good AI ideas is nearly infinite. Governance is required to prioritize those ideas into research that delivers the greatest impact within a cost and resource-constrained environment. Before launching into a project, Institutes should submit Research Proposals to the Board for their review and approval. Research Proposals should be specific and measurable, and document planned data and technology use, expected performance metrics, likely deliverables (e.g., code, tool, paper) and necessary infrastructure, cost, and resource considerations. Research Proposals should also include a plan for communicating findings to the public (to the greatest extent possible) and methods for sharing deliverables with other Institutes and the public as appropriate.

Despite exciting experimentation and a few small AI programs, the U.S. government is a long way from being “AI-ready³.” While a percentage of basic research must exist to explore and push technology beyond the realm of the possible, the Board’s review and approval of Research Proposals should lean towards applied research and focus on criteria such as technical feasibility, positive impact to the problem set, cost, schedule, and risk to accelerate our ability to be “AI-ready” now. A key aspect of the Boards review is to ensure effective coordination of effort between Institutes. For example, if one Institute is developing NLP techniques for entity extraction, another Institute can take advantage of that research and apply it to a different type of problem. The proposed AI-TAP (Goal 4) can be value add in enabling the board to accomplish this function. Suggested performance metrics for Goal 2 include:

- Number of Research Proposals submitted / reviewed / prioritized / funded
- Numbers of proposals that were able to leverage other projects work
- Number of proposal projects completed / shared with public
- Numbers / types of proposal outputs (data, software, publications, patents) delivered / shared (among Institutes, with the public)

Goal 3. Fund and Maintain the NAIRR Infrastructure. AI research requires significant amounts and types of data, which require significant storage and compute power. Noblis, in support of one of our clients for example, required storage capability for up to 10 petabytes of

³ NSCAI Final Report, pg. 2

data, processed data in excess of 10 terabytes, extracted, transformed, and loaded 1.5M records monthly, and accomplished processing and analytics of those data on a daily basis using AI to identify unique types of inconsistencies. While not all research projects may require this type of capability, large corpuses of data will need to be structured, stored, and analyzed. Synthetic data may need to be created where data does not exist or cannot be used due to legal or policy purposes. Problem-specific taxonomies will need to be created and testing and training data sets developed to account for wide-ranging variables and to address edge-cases. All the above are not trivial tasks and will require significant compute power to apply technologies such as NLP to automate data collection, processing, and synthesis so that research can be accomplished and delivered in a timely manner. These challenges and Noblis’ proposed recommendations will be discussed in greater detail in our response to question 1(E).

NAIRR, at the governance level should establish, maintain, and fund a hybrid cloud (e.g., AWS/AZURE) infrastructure that includes data storage and elastic compute and that makes tools available, as proposed in Table 2 (see 1 (D) response below) in support of the Institutes and their associated research. The NAIRR infrastructure should also support the AI TAP discussed in Goal 4. Suggested performance metrics for Goal 3 include:

- NAIRR hybrid cloud environment established / Cost to establish
- Number of Institutes with access to the infrastructure / Time to gain access
- Numbers / types of tools available / Numbers of new tools needed for specialized projects
- Numbers / types of data stored
- Numbers of compute requirements identified / funded
- Numbers / types of cloud and compute training required / funded
- Cost to maintain and use the cloud vs. cost to deliver AI output

Goal 4: Provide an AI Technology Accelerator Platform (TAP). Effective collaboration across the Institutes and more broadly across an AI Ecosystem requires an AI TAP. The AI TAP should provide a one-stop location for tracking and monitoring projects, for matching requirements to contract mechanisms, conducting evaluations, sharing best practices, gathering analytics, and presenting dashboards, all supported by security and access management protocols. Table 1 includes suggested features for an AI TAP.

Table 1. Noblis Proposed AI TAP Features and Benefits.

Features	Benefits
Project Setup	Establish new projects from an Idea Management Module
Evaluation	Capture evaluation findings in a user-friendly dialog window
Contract Management	Store contracts and process modifications at the project level for rapid traceability
Financial Management	Track invoicing data and principal investigator (PI) estimated expenditure data.
Milestone Tracking	Link milestones to dates with automated notifications; Alert PIs of upcoming deadlines
Document Upload	Enable PIs to load their deliverables and link them to milestones
Working Area	Provide secure document sharing and collaboration
Dashboard	Analyze and visualize project and portfolio performance indicators
Realtime Chat	Allow members to communicate in a secure collaboration environment
Security Authentication / Browser Encryption	Provide two-factor authentication with identify verification certificates and security tokens; Access a secure remote access portal using Secure Sockets Layer protocol
Role-based Access	Apply access based on user needs, as approved by the security official

Suggested performance metrics for Goal 4 include:

- Platform established; Levels / types of usage
- Numbers / types of projects; Submitted / Evaluated / Funded / Completed
- Numbers / types of users; Numbers / types of collaborations / Resulting in leveraged applications, new ideas, new partnerships, increased speed to delivery

Goal 5: Implement a Sustainable Acquisition Model. The NAIRR must create an acquisition model that allows customers (e.g., government) to fund research by performers (Institutes comprised on multi-disciplinary teams) in a highly efficient manner – which largely does not exist today. Amid a crowded federal marketplace, the acquisition model must provide valued services, speed, and ease of use on a single platform, such as what is being proposed via the AI TAP. Noblis proposes that the NAIRR initiate an AI Research Consortia that would operate and be supported by the AI TAP, and implement a decentralized, AI-focused acquisition approach to avoid redundancy, capitalize on best practices, and build a larger and more cohesive team of government and AI partner stakeholders. This type of approach would allow the NAIRR to be contract agnostic and identify, assess, and access a range of acquisition options, including Indefinite Delivery, Indefinite Quantity Task Order contracts, Government-Wide Acquisition Contracts, GSA Multiple Award Schedule contracts and non-FAR-based agreements, such as Other Transactional Agreements (OTA), providing for maximum scalability and accounting for likely fiscal constraints. OTAs in particular represent a useful vehicle as they focus on research, allow for exposure to large numbers of performers, and enable rapid capability transition. Also, the laws and regulations concerning intellectual property (IP) rights and cost accounting/pricing do not apply to OTAs although OTAs generally do contain provisions addressing such topics as IP rights and cost accounting. Suggested performance metrics for Goal 5 include:

- Numbers / types of contracts matched to Institute projects
- Time to establish a contract mechanism per project, Level of funding per contract
- Type of customer (e.g., government, venture capitalist) / Type of desired research or use

Goal 6: Provide Data Governance, Provide Policy Guidance and Recommendations, and Communicate Plans and Success. As noted in Goal 3, substantial numbers and types of data will be required for research. Obtaining those data, unless publicly available, as well as storing, processing, sharing, and noting its use in public are all likely to require data governance. While the identification of needed data and the creation of data sharing agreements should remain the responsibility of the individual Institutes doing the research, the Board should establish and maintain a Data Governance Committee comprised of Chief Data Officers to review the agreements, ensure legal and policy considerations are accounted for, and ensure a coordinated data approach exists across the NAIRR and its associated Institutes.

The use of data and new technology are likely to have policy impacts much like the recent discussions on biometrics. These impacts may result in the need for existing policy to be revised or new policy written. While the Institutes will have a role in unearthing these types of policy requirements, Noblis proposes that the Board establish and maintain a Policy Review Committee that routinely interacts with the Institutes to understand research-based policy implications and that proactively works with the executive and legislative branches of government to affect policy change that advances data and technology use, while ensuring proper protections to citizens.

A recent global consumer study conducted by Pega revealed that *many consumers couldn't even recognize some of AI's most basic tenets [and] that nearly half don't understand that AI*

*solutions enable machines to learn new things, and even fewer don't know it can solve problems or understand speech*⁴. In Noblis' experience, engaging with end-users (i.e., consumers) is critically important to enabling adoption, gaining support, and garnering important funding. For this reason, Noblis proposes that the Board establish a Strategic Engagement and Communications Team. This team would maintain a website that provides transparent information on Institute and NAIRR projects and successes, conduct outreach events to engage with local communities to help explain and highlight the value of AI, assist the Institutes in conducting hackathons, competitions, training, and recruiting events, and aid in the preparation and delivery of quarterly and annual reports and briefings to Congress, the President, the administration, and the public as appropriate. Suggested performance metrics for Goal 6 include:

- Data Governance and Policy Review Committees established / Strategic Engagement and Communications Team in place
- Numbers / types of data requirements / Numbers of Data Sharing Agreements in place / Levels of sharing allowed (e.g., fully open, just among Institutes, etc.)
- Numbers / types of policy requirements / Numbers of proposed vs. accepted refinements to existing policy / Numbers / types of new policy required, drafted, approved
- Communications website established / Numbers and types of listings, comments provided, by segment / Numbers of public meetings held / Briefings provided
- Numbers / types of events held; Numbers / types of attendees; Outcomes produced
- Annual report provided / Other reports (i.e., Institute Research Proposals) posted

B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:

B.i. An appropriate agency or organization responsible for the implementation, deployment, and administration of the Research Resource.

In our experience, the agency or organization responsible for the implementation and administration of a substantial “new start” such as NAIRR, must have significant standing, political backing, and influence to garner broad based support and long-standing commitment among participants, funding backers, and likely AI consumers/beneficiaries. Noblis recommends that the NSF oversee the day-to-day activities of the NAIRR, that they appoint a full-time Director, and that they implement proposed Goals 1-6 discussed in our response to question 1(a). The NAIRR Director should report to the OSTP Director who can provide guidance and convene the Board for higher level oversight and decision-making. Conceptualizing and codifying a plan of action is critical to organizational success. For this purpose, Noblis proposes that the NAIRR Director work with the Board to develop and deliver a NAIRR Program Objectives and Milestones (PO&M) document within the first 60 days of assignment. The NAIRR PO&M should describe agreed-to goals, objectives, activities, milestones, performance metrics, and responsible and accountable officials (by name). While the administrative functions remain critically important for establishing a strong foundation in the early months of “startup”, the NAIRR PO&M should include substantive and tangible deliveries of applied research within 180 days that are tied to challenging use cases to generate quick wins and show value. Throughout the first 180 days and for the remainder of the first year, NAIRR leadership should communicate with the President, Congress, the Administration, the Public, and with likely consumers (and

⁴ <https://www.ciosummits.com/what-consumers-really-think-about-ai.pdf>

funders) in no less than 45-day increments to report status, success, challenges, actions, and other pertinent information. Noblis proposes that the NAIRR PO&M be revisited on a no less than annual basis to discuss performance and needed refinements.

B.ii. A governance structure for the Research Resource, including oversight and decision-making authorities.

Please see our response to question 1(a), Goal 1: **Establish NAIRR Governance.**

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.

Noblis recommends that an AI TAP be implemented, and that hybrid cloud storage and compute capabilities be made available for effective NAIRR operations. Noblis employs a “lab to mission” methodology that follows three key principles: (1) validate and iterate; (2) simplify and uncomplicate; and (3) develop once, use many. As a company, we are tool and solution agnostic and use DevSecOps approaches that apply numerous capabilities and a wide variety of tools to deliver innovative research and solutions. Table 2 provides a list of frequently leveraged capabilities and is recommended as an initial set of technologies for consideration by the NAIRR and the Institutes. Noblis applies human-centered design and UI/UX processes to transform a stated need into a solution by engaging and collaborating with users/stakeholders throughout the design, development, and transition phases. Our process involves conducting evaluations on the iterative design (e.g., early mockups, prototypes, alpha/beta modules) and follows industry standards and best practices to ensure our tools allow users to accomplish their intended task. Noblis recommends that similar approaches be employed for NAIRR research efforts.

Table 2. Key Functions and Associated Tools Used by Noblis to Deliver Innovative Research.

Capabilities	Tools	Capabilities	Tools
Cloud storage	AWS S3	Multi-platform analytics and visualization	Grafana
High performance object storage	MinIO	Deep neural network training and inference	Tensorflow
Database (DB) management	Apache (Cassandra)	Multi-tasking operating systems	UNIX
Relational DB Management	PostgreSQL	Programming	Python
Relational DB construction	AWS RDS	ML Library	PyTorch
Large DB query	AWS Athena	Data stream processing	Kafka
Agile workflow	Atlassian Jira Confluence	Data integration	AWS Glue
Tester/Developer project management	VersionOne	Combine and query data across varied storage	AWS RedShift
APIs and SDKs for adding cognitive intel to apps	Azure Cognitive Services	Statistical analytics	SAS
Server scripting for dynamic web content	Node Js	Statistical computing	R
Number factoring	GNFS	Data clustering, classification, and regression	SciKit
Store, search, and compute (at scale)	Elastic	Automated code quality inspection	SonarQube

Capabilities	Tools	Capabilities	Tools
Networking for large scale data computation	Hadoop	Big data search and access	PowerQuery
Web & micro-service development & integration	Java Spring Boot	Big data (ML and Statistical) analysis	H2O
Automated software development	Jenkins	Data integration and analysis	Splunk
Software development via crowdsourcing	Mechanical Turk	Data visualization and analytics	Tableau
Container building and deployment	Docker	Data visualization	Kibana
Container clustering for app deployment/scaling	Kubernetes	Automated data extract, transform and load	NiFi
Container application management	Portainer	Probabilistic data structure	Bloom Filter

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource.

Noblis has extensive lifecycle data experience that spans collection, cleaning, transformation, and analysis involving a vast array of data types. Examples of government and commercial data that Noblis has worked with include web traffic, cell records, digital video, hard drive and removable media, social media, healthcare and geospatial data, basic safety messages and vehicle trajectory data, taxpayer, facility, financial, and air traffic information, and biometrics. Based upon our experience, too much of a data scientists time is spent discovering and preparing data for analysis vs. time spent on analysis to deliver valuable insights. The volume, velocity, and variety of data available today contains incredible value and can be applied to address myriad issues. Those same traits however make data search, access, retrieval, use, and analysis overwhelming. As part of the NAIRR governance process, and in accordance with relevant data sharing agreements, Noblis recommends that a thorough review be conducted prior to collecting, procuring, or distributing any datasets to the research community. Data acquired or generated should be specifically tied to use cases defined as high priority by the Board to avoid the risk of bad⁵ data entering the system. Table 3 lists challenges and recommended approaches to the use of high-quality open source and government data for NAIRR goals and objectives.

Table 3. Challenges and Recommended Approaches to the Use of High-Quality Open Source and Government Data for NAIRR Goals and Objectives.

Challenge	Recommendation
Data Structure: Estimates say that just 20% of data is structured, while unstructured data accounts for 80-90% ⁶ . AI/ML or any type of computing is best accomplished when data is structured and can be compared/analyzed within a relational database.	Leverage NLP to auto-curate and structure data that is not yet structured. Create mission or topic specific taxonomies that are continuously updated via automated entity mapping. Use training datasets for NLP of tagging pipelines such as Parts of Speech, topic modeling, and Named Entity Recognition.
Data Sensitivity (i.e., Personally Identifiable Information (PII) and / or classified government data): The escalation of security breaches involving PII has contributed to the loss of millions of records over the	Create and use synthetic data coupled with real data that can be validated to not bias the data samples. Incorporate API toggles that leverage multiple registration tags and that use AI to teach the system the range of available enrichments while also structuring the

⁵ Bad data may include data that is poorly structured, sensitive, bias, etc.

⁶ <https://monkeylearn.com/blog/structured-data-vs-unstructured-data>

Challenge	Recommendation
<p>past few years⁷. While large amounts of information are made available via social media, unauthorized use of those data for studies can have legal ramifications. AI/ML research often requires access to sensitive information that may contain PII or in other ways be deemed sensitive. If the sensitive fields are removed, the remaining data often has limited utility.</p>	<p>APIs so that they limit or restrict queries based on role-based access. Conduct pre-correlations between data sets to enable research. Removing specific addresses but leveraging zip codes from medical records for example can still provide information about economic and medical demographics of the community.</p>
<p>Data Stagnation, Corrupt Data, and Incomplete Large Data Sets: Data sets will need to continuously evolve in both scope (e.g., number of records) and breadth (e.g., number of fields) to avoid research being overcome by new data before it's even released. Data are often corrupted and simply removing observations that are not perfect may bias results.</p>	<p>Leverage NLP to auto-curate and structure data. Create mission or topic specific taxonomies that are continuously updated via automated entity mapping to reduce the need for researchers to formulate their own taxonomies prior to working on new algorithms.</p>
<p>Data Value: Enormous amounts of data are being produced daily. Most of it however is not mission or topic relevant and many sources (e.g., algorithms, bots, influencers) are not trustworthy, explainable, or credible.</p>	<p>Create a data rating schema to assess mission and topic value. Incorporate crowd-sourcing software and sentiment analysis to deliver credible and relevant data. Leverage deep and recurrent neural networks to create detection models that extract features and index and produce objects and relationships. Implement clustering models that generate terms of interest, import custom entities, and create customizable graphs.</p>
<p>Data Bias in Training Data Sets: Biases may be deliberate or inadvertent and may be reflected in techniques used to collect, generate, normalize, update, or analyze.</p>	<p>Leverage statistical techniques to validate that bias are not present. Incorporate and consistently apply a Data Quality Assessment to ensure that approaches to data collection ensure appropriate representation.</p>
<p>Data Robustness: Often when faced with AI development challenges, researchers will use a portion of the available data as a training data set, and the remaining data as a testing or validation data set.</p>	<p>Clearly distinguish between training data sets and testing data sets which may be specifically sampled to include deliberate attempts to stress algorithmic approaches to ensure they are robust to edge cases.</p>

H. A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector.

Noblis proposes that the NAIRR consider employing a model like the highly successful Fraunhofer-Gesellschaft (Fraunhofer Society)⁸ network of institutes for applied research. The NAIRR funding sustainment model, much like that of the Fraunhofer should include funding derived from diverse sources, including federal, state, and public funding fees. In Noblis' experience, obtaining government provided funding, allocated through multi-year Congressional appropriations are the most consistent form of funding. Other potential sources of funding to consider include those gained from Foundations such as the McCain Foundation, Bill and Melinda Gates Foundation, and the Warren Buffett Foundation, and through engagement with Venture Capitalist (VC). Critically important to ensuring industry and venture capitalist involvement is providing maximum flexibility for maintaining intellectual property (IP) rights, providing licensing fees for continued development and deployment of unique technologies and solutions, and encouraging spin-off companies – which in turn will add more expertise, development, and competition in the market. Allowances for maintaining and / or continued funding of IP will encourage private industry and VC investment. NAIRR funding should be provided via contract mechanisms, such as those discussed in Goal 5, pg. 5. Finding the proper

⁷ <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-122.pdf>

⁸ <https://www.nap.edu/read/18448/chapter/13#225>

contract vehicles and / or establishing effective contract mechanisms can be burdensome which is why we recommend the AI-TAP (Goal 4) as a means to rapidly match performers to customers via contract mechanisms and deliver applied research.

Customer involvement in research is critical to ensuring understandable, usable, and adoptable AI. For this reason, Noblis recommends that government (local, state, federal) be intimately involved in the creation of technology innovation pilots [based on well-defined and approved use cases – see 1(e), pg. 7-8] and that they and the Research teams routinely interface throughout the project, akin to the process described in 1(d), pg. 7. This type of involvement will generate quick wins, provide greater interest and value-add, and result in additional funding for projects.

*We should embrace the AI competition. Competition already infuses the quests for data, computing power, and the holy grail: the rare talent to make AI breakthroughs*⁹. Engagement with academia, specifically with students interested in AI technologies, in a way that implements a “dual system” of education/apprenticeship, should be a particular focus for NAIRR. As part of their undergraduate or graduate curriculum, students often must, or seek internships. This “free” or relatively inexpensive labor force infuses perspectives, grows talent, and results in opportunities for future employment. Noblis employs this type of “dual system” and is a proud recipient of WayUp’s Top 100 Internship Program award in 2020 and 2021. Noblis’ Internship Program gives undergraduate, Masters, and PhD students an opportunity to support direct client work and contribute to our internal research and development programs. This year, our internship program received nearly 7,500 applications for our 75 projects. Our interns came from over 50 colleges and universities and 18 states. In 2020, Noblis had an 80% acceptance rate on our intern-to-full-time offers. The opportunity for private industry to gain access to this type of talent and potentially hire home-grown experts will drive industry engagement and funding for NAIRR efforts. Noblis’ experience in establishing state of the art labs and provisioning them with a vast array of capabilities, tools, data, and expertise generates quality research and appeals to prospective interns and applicants alike. Noblis proposes that the Institutes have and maintain similar capabilities and tools as suggested in Goals 3 and 4 and in question 1(D).

We look forward to the opportunity to discuss our knowledge, skills, and capabilities to help the Task Force make progress in furthering the value of AI!

⁹ NSCAI, pg. 2

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Northeastern University

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

OSTP and NSF Implementation Plan for a National Artificial Intelligence Research Resource (NAIRR)

Northeastern University Response

Organization Information:

Name/Address of Organization: Northeastern University, 360 Huntington Ave, Boston, MA 02120

Business Size of Organization: 6,307 employees (as of Fall 2020)

Socio-economic status of Organization: 501(c)3

Northeastern Respondents: Jennifer Dy, Deniz Erdogmus, Usama Fayyad, David Kaeli, Ningfang Mi, Dana DeBari, Timothy Leshan

Northeastern University Responses:

1) What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list below) for which you are identifying options.]

The performance of AI and machine learning (ML) algorithms are highly dependent on the quality and size of available data. Training AI/ML models for large data requires higher computational and storage resources. Moreover, modern powerful ML models (deep learning, large language models, high-dimensional and sparse data sets) are increasingly more complex, requiring an increasing number of learnable parameters and repeated iterations – often requiring “bursty usage” of elastic computing grids.

The availability of massive data and computational resources is crucial in making progress in AI/ML, which is mostly available only to large private companies and national labs. Component D is one of the most critical components to have for a National AI Research Resource (NAIRR). Another important component is H to continue to support research in AI through federal grants in partnership with the private sector and national labs. See responses below for details.

Component D is also potentially very relevant to addressing two of the biggest bottlenecks facing any organization that wants access to cloud elastic computing:

- a. The difficulty of setting up a cloud environment (see further discussion below)
- b. The expense for data egress costs once a public or private cloud environment is up and running – while the cost of storage is cheap, the cost of moving clouds or moving data to other computational environments becomes prohibitive. Creating a shared storage cluster by NAIRR that allows high-speed low-cost transfers of data out of clouds would address a big issue and would create more competition between the big cloud providers.

Component F on data security is gaining prominence, especially when sharing data on a limited basis. Many organizations in finance, healthcare, and telecom are learning that the best way to “share data” without data leakage to potentially malicious organizations and nation states is by creating environments where the participants bring their applications to the data rather than bring data to the applications in the standard typical model. A compute cloud where data is stored centrally, and users are enabled to run their algorithms on the data without moving it (via data lakes or storage clusters coupled with significant compute) creates a more tenable approach to controlling data movement and unauthorized distribution of data. NAIRR can play a central role to enable such a model to ensure that data is only used by authorized parties and no indirect paths to share “copies”

are allowed. This also creates strong opportunities to impose privacy-preserving mechanisms (e.g., access via differential privacy) to further guarantee responsible use (Items G and F).

2) Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Over the past half-century, the United States has made significant investments in high-performance computing (HPC). These centers have designed powerful computing infrastructures that serve the computing needs of thousands of scientists and engineers as they generate new discoveries and remedies, enabling us to develop a better understanding of our world.

In 2010, high-performance computing (HPC) researchers from Lawrence Berkeley National Laboratory observed that cloud computing and data analytics were starting to dominate the data center, and so should be evaluated as a potential platform for HPC [Guida2020]. The study found that many scientific applications could be effectively scaled on cloud-based HPC platforms, but major impediments remained to fully leverage the Cloud for HPC applications, including memory system performance and communication costs. In 2015, Reed and Dongarra argued that the future of scientific discovery would combine advances from both machine learning and traditional HPC, producing a new and novel high-performance computing ecosystem [Reed2015]. Due to major differences in the characteristics of these workloads, it becomes ever more challenging to deliver a single computing framework. Today, many machine learning workloads are designed to leverage containerized services [Wang2018, Jouppi2020, Yi2020] and cloud-based hardware/software stacks [Cusumano2019] to process the growing volumes of data.

Elastic Cloud technology can provide an effective model that can deal with the heterogeneity issues associated with emerging HPC and machine learning applications. One major motivation for considering co-hosting these two classes of workloads is that many HPC applications generate mountains of data that need to be explored [Buffat2014, Balaprakash2019]. Further, HPC applications are beginning to leverage machine learning algorithms to efficiently steer simulations and iterative applications, leading to solutions in a fraction of the time [Wozniak2018, Kurth2018, Dong2020]. We believe that this emerging model will play an important part as we develop a national infrastructure to support exploration in artificial intelligence. This can be thought of as the next generation adaptive computation for delivering optimized AI workloads platforms – reducing time, cutting energy usage for sustainability, and maximizing throughput across more workloads.

There is a growing number of machine learning applications that require large-scale data access and processing. Scale-out storage infrastructure has become essential for storage systems (i.e., hyperscaler [Lu2018] and cloud-storage [Luo2020]) to provide vast space and high throughput. One promising direction of solutions is to deploy a disaggregated model on cloud-based ML platforms where storage drives are physically separate from the compute nodes. In such a system, compute and storage resources can be scaled independently for different needs, and resource management becomes more flexible. All-flash arrays have emerged in data centers recently, yielding superior performance to traditional storage devices. A hybrid of diverse storage devices, including traditional HDDs, SSDs with PCIe interface (e.g., NVMe), and phase-change memory devices (e.g., Intel Optane), can be installed in storage nodes, providing a variety of device capacities and processing speeds. Various scheduling schemes and management policies can further be applied on a single or a group of compute/storage nodes based on the demands of machine learning workloads. Another possible motivation for disaggregating the storage system is that the safety and integrity of (sensitive) data can be maintained by applying different data preservation policies on different storage nodes, which thus uncouples the access to private and public data repositories.

3) How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

As AI is widely deployed in a variety of applications affecting various sectors of society, we need to make sure AI systems are trustworthy: i.e., ethical, responsible, fair, reliable, secure, privacy preserving, safe, transparent, and interpretable. Research in trustworthy AI is in its infancy requiring convergence of experts from multiple disciplines, such as philosophy, law, sociology, psychology, and AI. To help advance this new field in AI, NAIRR through component H supporting federal funding in partnerships with the private sector can help accelerate development of this field. NAIRR can also help by establishing an appropriate agency that provides guidelines for ethical oversight. In designing a shared data and computing resource, NAIRR can help reinforce principles of ethical and responsible AI through components F (on security) and G (on privacy).

Furthermore, the approach to provide equitable access to computation and analytics is to work on lowering the bar for skills needed to access the technology. Leveraging NAIRR-class resources requires a high degree of technical skills, specialized talent and know-how, and the use is far from user-friendly. We need to reduce these barriers of entry and make these resources accessible to much larger groups of users by simplifying their use and by embedding pre-prepared packaged “solutions” to many of the common problems in analytics and Data Science. Usage can be increased by also making the outputs (e.g., analytics, insights, reports, etc.) much easier to produce and interpret. This simplification is much needed not only to increase the user base and make access more equitable, but to also reduce the significant redundant effort by advanced groups in building up the tools, utilities and environments to make their own work easier.

4) What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Two unique resources that are available at Northeastern are the DARPA Colosseum Network Simulator and the membership in the Massachusetts Green High Performance Computer Center (MGHPCC). Each of these facilities provide unique capabilities to the AI/ML research community.

NAIRR will need to acquire and provide AI-ready data sets. This is a new requirement that is needed by the AI/ML research community. Rich data sets from key domains (e.g., health, the exposome, 5G/6G communications, severe weather, the environment, etc.) will attract researchers from a broader range of areas to leverage the capabilities of NAIRR. We will need appropriate models for facilitating data use agreements, ensuring data security and privacy, while also considering the fairness and bias issues related to these data sets. With the major investment by Northeastern in the Institute for Experiential AI (focused on working with partners in its AI Solutions Factory and on Responsible AI practice), and other institutes in Robotics, Cybersecurity, Network Science, and Wireless IoT, Northeastern has many galvanized research and application resources dedicated to solving the critical problems to society.

In practice, much research work in ML that enables AI is evaluated and published over unrealistically simple data sets and schema. Real world data is much more complex: data schema and the variations in data quality and availability. Most such data sets are not available to the general research community. Thus, much of their work is of little relevance to real applications. However, there is no lack of data. The issue is the expense of gathering it, cataloging it, labelling it, etc. There is a tremendous amount of data available on the Web. However, crawling it, performing entity extraction, and then organizing and labelling it properly is prohibitively expensive. Delivering this service

through NAIRR will help to generate the next generation of challenge data sets and will go a long way in advancing research and in removing the obstacles that are perceived as insurmountable by researchers in Machine Learning and Data Science. Many such training sets can be also from simulations and outputs of the HPC applications, in addition to traditional web and social media sources. Many other complex and large publicly available data sets can be leveraged: financial markets, capital and equity markets, census data, environmental data, economic data, etc. The problem facing most researchers is getting this data ready for analysis, which is a huge challenge.

5) What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

At the Massachusetts Green High Performance Computer Center, we have aggressively explored public/private partnerships in both the initial buildout of our Center, and as a theme for ongoing research into public/private cloud services. Specifically, the Center was founded as a partnership between the State of Massachusetts and five Massachusetts universities (Boston University, Harvard, MIT, Northeastern and the University of Massachusetts system). Each partner contributed to build a state-of-the-art green computing facility that services researchers across the Commonwealth.

In addition to the creation of this facility, ongoing research at the MGHPCC has aggressively pursued unique public-private opportunities that leverage the Center. Two examples of this focus include the Massachusetts Open Cloud (MOC), an ongoing project that has created a self-sustaining at-scale public cloud based on the Open Cloud eXchange model [MOC]. The MOC serves as a marketplace for industry partners, as well as a service for researchers and industry to innovate and expose innovation to real users. The second example of a successful public-private engagement is the AI Jumpstart program. This state-funded project leverages the expertise in AI/ML of three leading academic institutions (Northeastern, Boston University, and Tufts) with small- and medium-sized businesses that are looking to leverage the benefits of AI/ML technology within their organizations [AI-Jumpstart]. The support from the state provides for acquisition of a large computing cluster specifically designed for the most challenging AI/ML workloads. The state funding supports initial engagement grants for the companies to use to work with the faculty and their students.

Scaling programs such as AI Jumpstart to the national level and creating a framework where many companies are extremely interested in innovating with AI and data, though are not able to because of lack of platforms and know-how, would be a huge contributor to the economy. Such a framework would enable academia and small and medium-sized companies, who have the talent and skills, to work with many companies and organizations who do not. This is akin to the building of railroads in the 1800's and the interstate highway system in the 1900's – both enabled tremendous economic expansion and opportunities – we need a similar infrastructure around data, AI and computation.

6) Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

For NAIRR to be successful in terms of meeting the needs of a broad spectrum of researchers and scientists, two major hurdles need to be addressed. The first is the cost of services, which the NAIRR framework should help to address in part. How do we continue to provide access to this infrastructure to everyone, and ensure that the facility can meet the needs of those with fewer financial resources? One model is to charge a fee to those institutions that have financial resources and utilize this fee to ensure there are ample resources reserved for those institutions that have limited funding. A fair and equitable financial model to accompany this facility would have to be developed and deployed, one that ensures its long-term vitality and sustainability.

A second barrier that needs to be overcome is to ensure that everyone receives a proper education on how to best use this infrastructure. This includes the effective use of middlewares and tools/libraries that support effective use of these resources by non-specialists. In terms of hosting this infrastructure, expenditures on the equipment and services should come with an equal amount of support to deliver training to the future users of this infrastructure. Academic institutions with strong programs in AI and Data Science seem well-poised to deliver these services.

References:

- [Guida2020]** Giulia Guidi, Marquita Ellis, Aydin Buluc, Katherine Yelick, and David Culler. 10 Years Later: Cloud Computing is Closing the Performance Gap, 2020.
- [Reed2015]** Daniel A. Reed and Jack Dongarra. Exascale Computing and Big Data. *Communications of the ACM*, 58(7), 2015.
- [Wang2018]** Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, and Kailash Gopalakrishnan. Training Deep Neural Networks with 8-Bit Floating Point Numbers. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, page 7686–7695, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [Jouppi 2018]** Jouppi, Norman P. and Yoon, Doe Hyun and Kurian, George and Li, Sheng and Patil, Nishant and Laudon, James and Young, Cliff and Patterson, David. A Domain-Specific Supercomputer for Training Deep Neural Networks. *Commun. ACM*, 63(7):67–78, June 2020.
- [Yi2020]** Yi, Xiaodong and Luo, Ziyue and Meng, Chen and Wang, Mengdi and Long, Guoping and Wu, Chuan and Yang, Jun and Lin, Wei. Fast Training of Deep Learning Models over Multiple GPUs. In *Proceedings of the 21st International Middleware Conference, Middleware '20*, page 105–118, New York, NY, USA, 2020. ACM.
- [Cusumano2019]** Michael A. Cusumano. The Cloud as an Innovation Platform for Software Development. *Communications of the ACM*, 62(10):20–22, September 2019.
- [Buffat2014]** Marc Buffat, Lionel Le Penven, and Anne Cadiou. High Performance computing and Big Data for turbulent transition analysis. <http://www.netlib.org/utk/people/JackDongarra/CCDSC-2014/talk15.pdf>, 2014. Online.
- [Balaprakash2019]** Prasanna Balaprakash, Romain Egele, Misha Salim, Stefan Wild, Venkatram Vishwanath, Fangfang Xia, Tom Brettin, and Rick Stevens. Scalable reinforcement-learning-based neural architecture search for cancer deep learning research. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–33, 2019.
- [Wozniak2018]** Justin M. Wozniak, Rajeev Jain, Prasanna Balaprakash, Jonathan Ozik, Nicholson T. Collier, John Bauer, Fangfang Xia, Thomas S. Brettin, Rick Stevens, Jamaludin Mohd-Yusof, Cristina Garcia- Cardona, Brian Van Essen, and Matthew Baughman. Candle/supervisor: a workflow framework for machine learning applied to cancer research. *BMC bioinformatics*, 19(18):491, 2018.
- [Kurth2018]** Thorsten Kurth, Sean Treichler, Joshua Romero, Mayur Mudigonda, Nathan Luehr, Everett Phillips, Ankur Mahesh, Michael Matheson, Jack Deslippe, Massimiliano Fatica, et al. Exascale deep learning for climate analytics. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 649–660. IEEE, 2018.
- [Dong2020]** Wenqian Dong, Zhen Xie, Gokcen Kestor, and Dong Li. Smart-pgsim: using neural network to accelerate AC-OPF power grid simulation. *arXiv preprint arXiv:2008.11827*, 2020.
- [MOC]** The Massachusetts Open Cloud, URL: <https://massopen.cloud/>
- [AI-Jumpstart]** Providing Massachusetts Businesses with an AI Jumpstart, URL: <https://innovation.masstech.org/AIJumpstart>
- [Lu2018]** X. Lu, J. Chiu, S.-J. Chao, and Y.-B. Ye, “Design of instruction analyzer with semantic-based loop unrolling mechanism in the hyperscalar architecture,” in *ICS*, 2018.
- [Luo2020]** S. Luo, G. Zhang, C. Wu, S. Khan, and K. Li, “Boafft: Distributed deduplication for big data storage in the cloud”, *IEEE Transactions on Cloud Computing*, vol. 8, pp. 1199–1211, 2020.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

NVIDIA

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



NVIDIA’s Response to the Office of Science and Technology Policy Request Regarding An Implementation Plan for a National Artificial Intelligence Research Resource

Federal Register Document: [2021-15654](#); Citation: 86 FR 39081
Submission Date: 1 October 2021

NVIDIA Contacts:

Jack C. Wells, Ph.D., Science Program Manager, NVIDIA
Ned Finkle, Vice-President, External Affairs, NVIDIA

Response to Question 1. What options should the Task Force consider for any of the roadmap elements A through I?

A. Roadmap Element A: Establish and sustain a National Artificial Intelligence Research Resource, and metrics for success.

The goal is to help extend and evolve US leadership in the development, deployment, and usage of information technology (IT) and curated data resources to support artificial intelligence (AI) research and application. In doing so, we will advance both economic and social well-being as society enters a new era in the information age, commonly referred to as the AI age. We believe new thinking should be applied to the existing federally owned or sponsored high-end computing facilities to transition these resources into the AI age. Data is considered the “fuel of AI”, lending to the idea that the curation and availability of public research data and AI models must be transformed, and availability of pre-trained models for reuse is just as important as core data sets in democratizing access to AI capabilities. AI tools and methods require hands-on usage, but are evolving rapidly, increasing the need for infrastructure and expert training. These hardware and software resources will need to be updated frequently until the rate of change stabilizes in the AI ecosystem. Therefore, a National Artificial Intelligence Research Resource (NAIRR) needs to be broadly accessible to individuals with relevant experience and training. AI tools are being applied throughout science, social science, and humanities research and development (R&D). As these tools become ubiquitous in the modern world, AI resources, education, and job-training, should be provided so individuals are equipped to leverage them in an ethical and responsible way. Undergraduate education curricula must be broadened to include introductions to data science and AI so students can achieve literacy in core concepts and the responsible consumption of AI tools.

We recommend establishing a broad, national AI research infrastructure based on a hybrid cloud strategy by integrating “on premises” AI resources, e.g., like those HPC facilities currently provisioned at federally sponsored entities, with the public cloud. NAIRR should leverage existing capabilities on premises at national facilities today, and these capabilities may need to be modernized and/or updated to support complex, data-intensive AI workloads. We recommend that the NAIRR Task Force perform an assessment of data centers to which NAIRR may have access, e.g., at universities, national laboratories, or other facilities to determine the capacity and readiness for NAIRR to potentially leverage. The Task Force should assess the need to construct new, energy-efficient data centers to support the NAIRR mission. To our knowledge, such a national assessment of datacenter capacity for research resources does not exist, and this is needed now.

NAIRR leadership needs to establish clear performance metrics for the success of a National AI Research Resource. For the priority vision of democratizing access to AI with respect to diversity, equity, and inclusion, NAIRR can demonstrate leadership by identifying clear performance targets once an established baseline is determined. Assessing the relevance of these and other performance metrics should be a regular activity for the NAIRR leadership team. Such metrics should include, but are not limited to, the following:

- Measures of data from curated datasets and pre-trained AI models, created and available in a manner to minimize bias in AI.
- The number of unique users per year, and the number of new, first-time users per year.
- User participation (i) by geography on a per population basis, (ii) by institutional type affiliation, (e.g., industry, university, national laboratory, or government agency).
- User participation by groups traditionally underrepresented in AI, or more generally science, technology, engineering, and math (STEM) disciplines.
- The number of students instructed per year.
- The number of training and education programs completed per year.
- The number of scheduled available compute hours, AI model throughput, throughput per unit electrical power expended, and total electricity use.
- The number and severity of data-security events and attacks on NAIRR.
- The number and quality of scientific publications, patents, and achievements highlights published per year.
- The number of small and medium-sized companies participating and using NAIRR.

- B. Roadmap Element B. Create a plan for ownership and administration of the NAIRR. This includes identifying (i) an appropriate agency or organization responsible for the implementation and deployment of the Research Resource; and (ii) a governance structure for the Research Resource, including oversight and decision-making authorities.

It is our judgement that a federation of new and existing resources at federal agencies and other entities is the preferred approach for implementation of NAIRR, as each agency will understand best their mission needs, research community, and data resources. There is no single agency or entity that currently possesses the necessary capabilities for NAIRR mission success. (See subsection D below for discussion of necessary NAIRR capabilities). All agencies and private entities contributing to NAIRR will need to develop internal strategies and action plans to build the requisite capabilities to sustain national leadership in AI.

We recommend a single NAIRR program office be created to provide program ownership, oversight, and administration for each of the federated agencies to ensure that the federated model runs smoothly and so each agency provides the appropriate level of support to achieve the objective of overall national AI leadership. Having one entity with oversight responsibility is imperative for NAIRR to operate successfully and for the needs of the nation to be met in a timely fashion. We believe this NAIRR office should reside within the Commerce Department, possibly within the National Institute for Standards and Technology (NIST), with dotted line connection to each of the participating agencies. (See Roadmap Element I below for brief discussion of possible agency roles.) The Commerce Department plays the key role in the National AI Initiative Act (NAII) of 2020, e.g., in establishing the National AI Advisory Committee (NAIAC) tasked with advising the President and the National AI Initiative Office on topics related to the National AI Initiative.

- C. Roadmap Element C. Develop and implement a model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources.

NAIRR should be a federally governed and sponsored resource available for external use to advance AI research, scientific and technical knowledge, education, and workforce training:

- NAIRR should be open to all interested U.S. based users with appropriate institutional affiliation (government, private industry, higher education, national laboratories, non-profit institutions). Requirement for institutional affiliation facilitates risk management.
- User programs need to have well-defined goals and manage resources in an equitable and transparent manner.
- User fees should not be charged for non-proprietary research and if the user intends to publish the results in the open literature. NAIRR should provide sufficient resources for users to conduct work safely, ethically, and efficiently. Proprietary users are obliged to pay full-cost recovery.

Allocation mechanisms need to be deployed in alignment with National AI Initiative (NAII) goals. For programs focused on large resource consumption, a rigorous merit review of proposed work is required. For research projects funded by federal agencies, NAIRR resources may be requested through the corresponding agency sponsor program manager. For users requiring small resources, an open queuing system could suffice.

We recommend decision-making authority resides with a NAIRR Director, ideally someone who is a member of the research community and recognized for achievements in data curation and sharing, AI, HPC, and leadership of high-end HPC facilities. This director should be appointed by the Secretary of Commerce through consultation with the Director of the National Artificial Intelligence Initiative (NAII) Office. The establishment of a NAIRR Board of Directors can provide oversight on strategy, budget proposals, partnerships, and executive level hiring. We suggest the NAIRR Director appoint a Computing Program Manager and Data Program Manager to integrate and manage computing and data as interdependent but distinct programs. Additionally, two separate advisory boards – one for the social sciences and humanities community and another for the physical, biological, and

computer sciences and engineering communities, shall be established to represent their needs, requirements, and views. To engage the user community and facilitate information sharing, we recommend creating a NAIRR User Group Executive Committee. Members shall be elected to serve staggered three-year terms to preserve and propagate effectively institutional knowledge. Members should elect the User Group Chair and Vice Chair.

- D. Roadmap Element D. Establish capabilities to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.

There are at least five main capabilities required for effective operation of NAIRR: (a) a trusted steward and broker for the curation of complex public and private data sets and pre-trained, trustworthy, privacy-preserving AI models; (b) expert deployment, effective operation, and management of federated computing and data infrastructure; (c) delivery of educational tools and programs in advanced AI methods at the appropriate level (primary, secondary, undergraduate); (d) expertise in the operation of user support programs for research and education; (e) excellence in cybersecurity and AI operations, including a well-balanced security posture and risk-based vulnerability management across diverse program requirements from open to protected.

(a) Steward Data Sets and AI Models: Capabilities to serve curated datasets and pre-trained models owned by both government and non-government entities are required to extend US leadership into the AI era. It is the availability of large, curated datasets, like ImageNet¹, that added an essential element launching the current revival of AI research.² Known "reference" AI models can provide an effective floor of "best in class" to measure against and ensure an increasing level of rigor within specific domains. The primary shift in the new HPC-AI era is the importance of data integration with HPC facilities. These data resources need to be planned, curated, and made available for multiple domains using FAIR principles³ so researchers may develop and validate new AI models and methods. NAIRR should consider the merits of national standards for AI information exchange and the creation of AI tools, (e.g., recommender systems, continuous integration and testing technologies) to facilitate the quality, discovery, and sharing of research data and AI models.

(b) Operate Federated Data and Computing Infrastructure: As described in Roadmap Element A, a hybrid-cloud strategy proves to be the most efficient approach to AI infrastructure deployment. NAIRR should deploy a diversity of production-ready, proven platforms for AI, including necessary software and hardware. NAIRR platform testbeds demonstrate readiness of emerging AI technologies for future deployments. Federal agencies, e.g., National Science Foundation (NSF), Department of Energy (DOE), and Department of Defense (DOD), have expertise in managing HPC platform testbeds to advance their missions. (See Roadmap Element I.)

¹ <https://www.image-net.org/>

² Li, Fei-Fei, [How we're teaching computers to understand pictures](#), retrieved 16 December 2018

(c) Deliver Educational Tools and Programs: Given the ubiquitous nature of AI, and to reinforce the U.S.'s competitive position in the technology sector, there is an imperative to incorporate AI elements both technical and non-technical into U.S. primary and secondary education (K-12) curriculum, to ensure students can thrive as AI consumers, creators, professionals, and citizens. Youth need to understand how AI algorithms make predictions that impact their experience and decision-making, how this knowledge will ensure their rights are protected, and how working with AI is a viable career path. Unlike computer science (CS) education, AI education includes societal elements that can be taught in earlier years and should be woven across several curricula, including math, science, humanities. The path for NAIRR to incorporate broad AI curriculum has a strong leg up, stemming from the body of work done by several evidence-based organizations as well as after school and summer programs. Examples include AI4K12, AI4All, The AI Education Project, and the Boys and Girls Club of Western Pennsylvania's AI Pathways Toolkit. (Disclosure: NVIDIA has in the past or is currently funding these projects.) (See also Roadmap Element G.) Such programs will contribute to a national pipeline for AI literacy and ensure jobs can be filled as innovation grows.

(d) Operate User Support Programs: User support and experience is critically important and should be offered as tiers for job training, education, and research. For example, today's US federal supercomputing centers provide state-of-the-art HPC user support, such as user portals, expert liaison collaboration, and hackathon-style group learning experiences. NAIRR user support and outreach professionals must represent the diversity of our nation to advance equity and inclusion vision. (See performance metrics in Roadmap Element A.)

(e) Excellence in Cybersecurity and AI Operations: Strong capabilities in cybersecurity are necessary for the safe and secure operation of NAIRR. Given the prominent place NAIRR will play in our nation's innovation ecosystem, NAIRR will be a priority target for malicious actors. First-class capabilities for segregation of infrastructure control layers and user-control layers to enable "zero-trust" architectures⁴ should be deployed. Federal and private-sector investment to make AI models themselves robust, not just the infrastructure running AI⁵. While the AI model may be authenticated and protected from tampering, creating models that are robust to adversarial attacks requires investment in research and practice. (See also Roadmap Element F.)

E. Roadmap Element E. Implement an assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource.

Open access to vast sources of data available to the scientific and governmental enterprises and curated according to FAIR principles will promote development of new AI capabilities and enable AI researchers to focus on problem-solving. "Curating data in its place"

⁴ S. Rose, O. Borchert, S. Mitchell, S. Connelly. Zero Trust Architecture. NIST Special Pub. 800-207. <https://doi.org/10.6028/NIST.SP.800-207>

encourages public trust in the data quality and is consistent with the recommended federated approach to NAIRR management.

F. Roadmap Element F. Design an assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls.

In securing AI supply chains for most operating environments, it's helpful to consider datasets as code and models as binaries in evaluating control environments for AI workflows. Strong storage and provenance controls are needed to ensure no erroneous or malicious data is introduced via supply chain or data poisoning attacks^{6,7}. Operational capabilities that reproduce and explain models back to primary data becomes critical in incident response as well as advancing objectives for ethics, bias, safety, and other mission outcomes. (See Roadmap Element G). Google's "[Model Card](#)" proposal is one example how to organize the essential facts of machine learning models in a structured way. As individual AI models are composed into complex systems, model control through attestation and signatures is required throughout the ecosystem to increase confidence and trust in deployed models.

To secure deployed AI training pipelines and inference models, traditional security frameworks for core infrastructure and edge deployments like [AICPA SOC-II](#) and [NIST 8320](#) should be the norm, with additional emphasis on zero trust service design⁴ and strong data-at-rest controls to ensure multi-tenant isolation. Notably, model extraction⁸ and membership inference attacks⁹ may leak information about Intellectual Property (IP) and data used to generate a model, or the model itself, leading to IP loss and potential impact on privacy preservation within deployed models. Likewise, adversarial attacks on model decisions, and AI defenses may be used to evade or influence expected behavior of composed systems. New observations unique to AI derived from efforts such as the [Mitre AI Attack Framework](#) by leading entities like Microsoft, NVIDIA, IBM and others, should be integrated to further refine operational security for AI models.

Vigilance is required to sustain AI security. Vulnerability management and robustness will inevitably change over time as target markets change and novel attack patterns are developed. AI operations must continuously evaluate and update AI models in the field to ensure the overall health of the ecosystem. (See Roadmap Element G.) Active vulnerability management like [FIRST/Mitre](#) should be developed to enable operations teams to effectively manage risk of deployed models. Emerging technologies in confidential compute should be evaluated in testbeds to further strengthen isolation guarantees in shared environments. Standards developed by Trusted Computing Group ([TCG](#)) and Confidential Computing Consortium ([CCC](#)) that unify system level security, trusted enclaves, and operational software layers provide reasonable roadmaps for engagement

⁶ <https://arxiv.org/pdf/2005.00191.pdf>; <https://arxiv.org/pdf/1905.13409.pdf>.

⁷ <https://www.fireeye.com/blog/threat-research/2020/12/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor.html>.

⁸ <https://arxiv.org/pdf/1909.01838.pdf>; <https://arxiv.org/pdf/1806.05476.pdf>.

⁹ <https://arxiv.org/pdf/2007.14321.pdf>.

- G. Roadmap Element G. Deploy an assessment of privacy and civil rights and civil liberties requirements associated with the NAIRR and its research. (See also Question 2.)

Societal elements of inclusion and equity are important to address in AI education and research, especially given the potential negative aspects of AI in policing, criminal justice, financial systems, and housing discrimination, which are more likely to impact under-represented communities. The significant lack of representation in the tech industry of people who experience and can work to address the potential harms of AI applications only exacerbates this problem, so it is critically important that AI education is available to communities that have been traditionally underrepresented in technology. The task force should consider whether certain uses of AI systems should be prohibited or out-of-scope for NAIRR because they have the potential to violate privacy rights, civil rights, or civil liberties. Such uses include, but are not limited to, using facial recognition technology (FRT) to infer intent from facial expressions, using FRT to identify a person's sexual orientation¹⁰, and using FRT to identify a person's group affiliation¹¹.

Performance standards are needed for AI systems that adversely affect individual freedom. For example, AI systems that purport to predict a person's likelihood of committing a future violent crime are known to discriminate against certain groups. AI systems that do not meet a minimum accuracy should be prohibited from use by law enforcement and judges. Such models must be continuously revalidated over time to ensure they sustain predictive accuracy in the face of social, generational, evolving regulatory and other changes. Standards are needed for developing synthetic content and supporting the creation of tools for consumers to identify synthetic content.

- H. Roadmap Element H. A plan for sustaining the NAIRR, including Federal funding and partnerships with the private sector.

The predominate funding to establish and operate NAIRR will be federal funds. The private-sector financial investment in establishing NAIRR will be limited and focused, with participation likely to occur through provisioning of open software, procurements to provision AI platform infrastructure, and collaboration on grand-challenge research projects. NAIRR will be sustained through robust budget planning over a 5-year horizon. Long-term operational contracts should be established with universities, national laboratories, and federal agencies to host NAIRR infrastructure, as well as data services providers and public-cloud providers. NAIRR requires long-term access to modern datacenter infrastructure. (See Roadmap Elements A, D(d), and Question 2.) International cooperation with strategic allies, as appropriate, on the testing, evaluation, deployment, and sharing of resources for trustworthy AI systems, is also crucial to sustain NAIRR.

- I. Roadmap Element I. Implement parameters for the establishment and sustainment of the NAIRR, including agency roles and responsibilities and milestones.

¹⁰ Wang, Y., et al. (2018). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *J. Per.s Soc. Psychol.*, 114(2), 246–257. <https://doi.org/10.1037/pspa0000098>

¹¹ Kosinski, M. (2021). Facial recognition technology can expose political orientation from naturalistic facial images. *Sci Rep* 11, 100. <https://doi.org/10.1038/s41598-020-79310-1>

NAIRR should be led by a new program office within the Commerce Department, (as indicated in Roadmap Element B), with support from federal agencies and the Networking and Information Technology R&D (NITRD) Program, Interagency Working Group for AI R&D. NAIRR should publish a report on requirements for future resources every two years and an annual self-assessment report.

Other federal agencies have competencies and core capabilities that will contribute within the NAIRR federation. The National Institute of Standards and Technology (NIST) has competency in establishing trustworthy standards for data curation and AI, which is a critical need for establishing broad acceptance of AI solutions. NIST has a strong focus on cultivating trust in the design, development, use, and governance of AI systems by (i) conducting R&D to advance trustworthy AI technologies, (ii) establishing benchmarks, (iii) developing data and metrics to evaluate AI, (iv) developing technical AI standards, and (v) engaging in discussions toward development of AI policies. The National Science Foundation (NSF) possesses leading competency in administering traditional on-premises HPC research resources with availability of public-cloud research resources. The NSF has also demonstrated leadership in a variety of core and crosscutting programs, through facilitating access to cloud computing resources through the CloudBank program, and through a variety of educational programs at all levels – pre-K through colleges and universities. The US Department of Energy (DOE) and the Department of Defense (DOD) who support HPC, data, and high-speed networking facilities. The DOE owns national supercomputing and high-speed networking user facilities and curated scientific databases operated within the DOE complex of multiprogram national laboratories^{12,13}. The DOD's HPC Modernization Program delivers world-class high performance computational capabilities to the DOD's science and technology (S&T) and test and evaluation (T&E) communities. Other agencies, such as the National Institutes of Health (NIH), the National Oceanic and Atmospheric Administration (NOAA), and the National Aeronautics and Space Administration (NASA), also operate federally owned supercomputing data facilities and should be included within the NAIRR federation.

Response to Question 2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

NAIRR's highest and most urgent priority must be to establish itself as a trusted steward and broker for curated data sets and pre-trained AI models for use in broad areas of research and education. The complement to this is the exhibition of excellence in cybersecurity and AI operations.

Response to Question 3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and

¹² <https://science.osti.gov/User-Facilities/User-Facilities-at-a-Glance/ASCR>

gender equity, fairness, bias, civil rights, transparency, and accountability? (See also Roadmap Element G.)

- Require transparency of AI that can cause harm to humans, by prohibiting trade secret protections of AI used in law enforcement and criminal justice applications, particularly when individual freedom or equal access to resources are at stake. Adopt standards against which AI outcomes are measured. Robust transparency will include robust review, including accessibility to Freedom of Information Act (FOIA) requests.
- Bias in AI system outcomes often can be traced back to dataset inputs, in training or in production. Facilitate the creation of and access to large, clean datasets for R&D.
- Provide consistent, practical definitions of key principles, e.g., fairness, safety, accuracy, performance, etc. Create tools to assess and rate AI systems on such key principles. Require transparency through institutional review board (IRB) review and/or community stakeholder board (CSB) review/approval, as appropriate, of AI systems effect on humans.

Response to Question 4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

- Leading edge AI platforms from US-based accelerated-computing vendors, in partnership with universities, laboratories, and facilities. An example is the University of Florida-led partnership around the HiperGator AI Supercomputer, (<https://ai.ufl.edu/>).
- Widely adopted software frameworks¹⁴ supporting AI development and applications.
- World-class datacenter capabilities at NIST, NSF, DOE, NIH, and DOD HPC Centers.
- World-class, high-speed, wide-area networking (WAN) through, DOE's Energy Sciences Network (ESNet), and other national scale WAN resources.
- World-leading, US-based cloud service providers (CSP) of compute and data infrastructure.
- Public data from federally funded research promoting FAIR principles, e.g., at Data.gov, NIH's Open Domain-Specific Data Sharing Repositories¹⁵, DOE's PURE Data Initiative¹⁶.

Response to Question 5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public-private partnerships will enable NAIRR to start faster, deploying today's robust AI infrastructure platforms quickly. Federal agencies have long partnered with the private sectors in procuring and/or operating HPC research resources through management and operating (M&O) contracts. This approach has been effective in delivering world-leading capabilities by transferring selected risk management – and related risk premiums – to universities or private contractors. These represent positive exemplars available for NAIRR to model, and long-term M&O contracts should be established within the NAIRR federation as communicated in response to Roadmap Element H. Newer programs designed to make public-cloud research resources

¹⁴ <https://www.nvidia.com/en-us/ai-data-science/>

¹⁵ https://www.nlm.nih.gov/NIHbmic/domain_specific_repositories.html

available (e.g., CloudBank or Strides) indicate that long-term contracts are appropriate to establish with multiple cloud-service providers.

Significant international exemplars exist of AI research resources promoting public-private partnerships. The AI Bridging Cloud Infrastructure¹⁷ (ABCI) center in Japan is a designated AI supercomputer resource open to public and private research offering cloud access to compute and storage capacity for AI and data analytics workloads. ABCI's software environment is a container-based ecosystem with core programming and other tools as part of its standard offering. NVIDIA launched Cambridge-1¹⁸, the UK's most powerful supercomputer, enabling top scientists and healthcare experts to use the combination of AI and simulation to speed the digital biology revolution and bolster the UK's world-leading life sciences industry.

Response to Question 6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

- In the absence of a well-understood definition of success, NAIRR may be challenged over time to maintain focus on its mission. The vision to “democratize access to AI R&D” is both compelling and vague. Well defined success metrics are needed to sustain NAIRR, and, in the process broaden access of resources for the US research community.
- Growing the participation of traditionally underrepresented groups will be a challenge requiring full participation and coordination with many local, regional, and national programs. Some potential limitations include (i) lack of visibility into entire AI R&D landscape; (ii) marginalized and underrepresented communities not at the table to create the plan, and (iii) absence of an actionable pipeline framework to build AI R&D readiness. A comprehensive strategy should be created in which NAIRR has specific, coherent actions to achieve the program targets in these, but with underrepresented and historically marginalized communities included up front and contributing to that plan.
- The user programs and communities for NAIRR are not well defined or understood. NAIRR should create well-defined user programs with public outreach for participation to bring focus and clarity to the community.
- The US is limited in access to science data repositories provisioned to support AI. A high priority is to leverage existing science-data-generating programs to produce world-leading data repositories ready for use in building AI models. These data repositories should be created in collaboration with the subject-matter experts and in adequate proximity to HPC resources.

¹⁷<https://abci.ai/>

¹⁸<https://blogs.nvidia.com/blog/2021/07/07/ceo-unveils-cambridge-1/>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Open Commons Consortium at the Center for Computational Science Research, Inc.

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

86 FR 46278 Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource

Submitted by the Open Commons Consortium (OCC), a Division of the Center for Computational Science Research Inc. (CCSR), a 501(c)(3) not-for profit located in Chicago, IL.

Question 1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

Question 1D. D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

Response. We argue that such a shared computing infrastructure should be based upon a **data commons**, which is software platforms that co-locates: 1) data, 2) cloud-based computing infrastructure, and 3) commonly used software applications, tools and services to create a resource for managing, analyzing, integrating and sharing data with a community [1–3]. An example of a data commons is the NCI Genomic Data Commons [4] that is used by over 50,000 researchers each month and provides over 1 PB of data to the research community in an average month.

More information about data commons is contained in the Response to Question 4 below.

Question 1E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

Response. As in the response to question 1D above, we propose that government datasets supporting the National Artificial Intelligence Research Resource should be made available to the public using a data commons. More specifically, we propose that a data commons with open FAIR APIs [5] be used as the foundation for make the data available to the public.

Question 2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

Response. We argue that data commons developed to support different geographic regions (**regional data commons**) provide a good foundation for the NAIRR.

An example of a regional data commons is the Chicagoland COVID-19 Commons, which is an instance of the Pandemic Response Commons that is operated by the Open Commons Consortium. The Chicagoland COVID-19 Commons contains data COVID-19 related data from the Chicagoland and Illinois region, including case and fatality data, vaccination data, clinical

data from COVID-19 patients provided by regional healthcare providers, SARS-CoV-2 strain data, health disparities data, and related data.

Multiple regional commons can be integrated together to form a national data ecosystem [6] to tackle an AI problem of interest, while still reflecting important regional differences in the data, as well as regional differences in how can be best be analyzed to serve its region. Multiple regional commons can also be used as a foundation for federated machine learning.

Importantly, regional data commons can better reflect and engage with the local community, which provides a basis for reducing bias and increasing diversity of the data it supports.

Although it can take a while to set up a regional data commons, once it is set up with the appropriate consortium governance agreements, data governance agreements and commons governance agreements, the regional commons can be quickly repurposed to collect new data types, support new projects, and respond quickly to emergencies, such as providing a data driven foundation for new public health emergencies.

For this reason, prioritizing setting up regional data commons provides a good foundation for the NAIRR.

Question: 3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Answer. An important component of the OCC Pandemic Response Commons approach has been to engage with the regional community through a Community Engagement and Outreach Working Group. This approach is possible because we are developing and operating a regional data common with close ties to the local and regional community.

Question 4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Answer. Data commons developed and operated by the Open Commons Consortium are developing using **open source software** (Gen3, <https://gen3.org>); support **open and FAIR data** through Gen3's open APIs; support **reproducible research** through Gen3's use of containerizing workflows that access data with persistent opaque identifiers; and its consortium membership agreements specify that research results be published in **open access** journals whenever possible.

The Open Commons Consortium has standard and time-tested agreements for: i) Consortium membership and governance; ii) contributing data to commons (data contribution agreements); accessing and analyzing data from (data use agreements); iii) setting up and operating working groups around projects of interest; and related activities.

In short, data commons developed using the open source Gen3 software and operated by the not-for-profit Open Commons Consortium for consortia provide a good foundation for the NAIRR.

Question. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Answer. A good model might be to support multiple different types of public-private partnerships serving different roles in the NAIRR. As an example, CCSR supports the BloodPAC Consortium, a private-public partnership that was originally launched as part of the Cancer Moonshot that accelerates the development, validation and accessibility of liquid biopsy assays to improve the outcomes of patients with cancer. The BloodPAC Consortium is a consortium of over 50 member organizations, including universities, commercial companies, and USG agencies. The BloodPAC Consortium develops and operates the BloodPAC Data Commons to provide its members and the broader liquid biopsy community with a data driven approach to advance research through the open sharing of data and its analysis.

References

- 1 Grossman RL. Data Lakes, Clouds, and Commons: A Review of Platforms for Analyzing and Sharing Genomic Data. *Trends Genet* 2019;**35**:223–34. doi:10.1016/j.tig.2018.12.006
- 2 Grossman RL, Heath A, Murphy M, *et al.* A Case for Data Commons: Toward Data Science as a Service. *Comput Sci Eng* 2016;**18**:10–20. doi:10.1109/MCSE.2016.92
- 3 Heath AP, Ferretti V, Agrawal S, *et al.* The NCI Genomic Data Commons. *Nat Genet* 2021;:1–6. doi:10.1038/s41588-021-00791-5
- 4 Grossman RL, Heath AP, Ferretti V, *et al.* Toward a Shared Vision for Cancer Genomic Data. *N Engl J Med* 2016;**375**:1109–12. doi:10.1056/NEJMp1607591
- 5 Wilkinson MD, Dumontier M, Aalbersberg IJ, *et al.* The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* 2016;**3**:160018. doi:10.1038/sdata.2016.18
- 6 Grossman RL. Progress Toward Cancer Data Ecosystems. *Cancer J* 2018;**24**:126–30. doi:10.1097/PPO.0000000000000318

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Oracle America, Inc.

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document

October 1, 2021

Attn: Wendy Wigen, NCO
Office of Science and Technology Policy and the National Science Foundation
2415 Eisenhower Avenue
Alexandria, VA 22314
Submitted electronically via email: [REDACTED]

Dear Ms. Wigen and NAIRR Task Force Members:

On behalf of Oracle America, Inc., thank you for the opportunity to provide input on FR Doc. 2021-15660, "Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource" (RFI). [Oracle](#) is a global leader in data management solutions, information technologies and cloud computing, serving 430,000 companies across 175 countries, and supporting hundreds of research projects and millions of students in more than 100 countries through its outreach programs. From its beginning, Oracle has been committed to supporting and advancing research, investing \$56 billion in research and development over the past decade. Oracle appreciates the importance of developing a shared research infrastructure to provide Artificial Intelligence (AI) researchers and students across scientific disciplines with access to computational resources, high-quality data, educational tools, and user support, and fully supports this effort. We believe if the NAIRR initiative succeeds, it could be for AI what DARPA's sponsorship of Ethernet and Internet research was for the creation of the modern information age.

In the following pages, we will address the specific points of the RFI in more detail. At a high level, we believe the NAIRR Taskforce and the resulting resource should be grounded in five primary principles to achieve the outcomes set forth in the RFI and ensure the NAIRR resource is able to evolve as technology evolves:

1. A common, flexible, open software framework that supports easy exchange of and collaboration around data, AI model, and compute resources ("run anywhere");
2. Open standards for data and AI models;
3. Open standards for fairness, privacy and security in AI;
4. Open, sustainable architecture that avoids dominance by any single vendor or entity; and
5. Support for a diverse range of students, educators and researchers.

Question 1 - Oracle offers the following feedback on goals E-I:

Regarding Goal E - The dissemination and use of government data sets and security and access controls are fundamentally connected issues related to accessing, sharing and managing data in a shared infrastructure model. We believe NAIRR can and should play a critical role in disseminating data, providing infrastructure to facilitate use of that data, and protecting the rights and privacy of individuals that could be violated by the availability of data collected by the government and by corporate and academic entities.

Although the Federal government currently publishes vast quantities of data via organizations like NOAA and NASA, and resources like www.data.gov, inconsistent data management practices and

structures, data set size, compute costs, network bandwidth limitations, and inconsistent use licenses can make this data difficult and costly to understand, utilize, and combine with other data sets. Data sets that contain personally identifiable information (PII) present additional challenges: while HIPPA protects health data to some degree, similar controls are not available to protect the privacy of individuals in other large data collections, such as faces and license plates captured by street view mapping processes or video doorbell monitors. When researchers combine public and/or anonymized data sets with data sets containing PII, as for medical clinical trials, further use and privacy complications ensue.

We recommend a “data-as-a-service” (DaaS) approach from the outset that sets forth clear policies and practices related to access, security, metadata, search, retention, and use, with specified processes for updating these policies and practices over time as technologies and data evolve. While to date, DaaS mostly has been used to monetize data sets for business reasons, we propose something new – ***Data as a Service for Research – or DaaSR*** – implemented to enable wider use of more data for AI modeling and research, consistently and securely, and with built in safeguards to manage data ownership, data provenance, data security, and data governance.

A *DaaSR* approach to NAIRR could provide the following benefits for data set owners and users:

- Avoidance of vendor lock-in
- Standard metadata for data content, security, search and lifecycle
- Ability to support data in a wide range of formats
- Simplified search across data sets to enable more researchers to find and use the data they need
- Consistent data management and practices to enable combination of data sets and compliance with retention and reproducibility requirements
- Separation of the analysis/presentation layer from the actual data itself (this is the most open approach)
- Ease of administration, collaboration, and audit

To ensure that NAIRR meets its goals of being generally usable and widely accessible, in establishing a *DaaSR* approach, we recommend that NAIRR:

- engage with standards organizations to define open standards, a data taxonomy, metadata, and best practices for data to ensure interoperability, using existing standards, modalities, and taxonomies where possible,
- define and apply consistent and appropriately differentiated data licenses across all government-collected data,
- establish privacy and security standards and security policy profiles to govern data access and use in ways that balance individual privacy protections with advancing research interests,

ensure data is encrypted at rest as well as in transit to prevent any unintentional breaches, while including the encryption method and routines in the metadata, and

- establish clear and consistent rules for data retention and maintenance, and preservation and maintenance of the workflows and tools that enable use of that data.

Regarding Goal F - Over time, more data will be made available, in differing formats (text, graph, video, audio, etc.), and more privacy concerns and protections will be introduced. In short: the challenges with publishing, using and protecting data are only going to grow. The existence of a consolidated research data resource like the NAIRR will amplify the already material risks that bad actors may try to access and use data for nefarious purposes, and the expanded application of AI contributes to a growing risk that individual privacy and autonomy may be compromised, even accidentally, by collation of apparently innocuous, anonymized, individual data points.

To ensure the successful implementation, security and sustainability of the NAIRR, these challenges must be addressed strategically. Using a *DaaS* model, access to data sets is enabled according to graduated security policies that reflect the risks associated with each data set, while data can be located anywhere and searched from anywhere. In such a model, data is “fluid” but controlled, enabling data owners to control cost, access to and retention of their data sets within an agreed set of guidelines and policies defined by NAIRR. These guidelines and policies should include a defined data catalog structure and consistent metadata to tag and define data sets. This overarching policy and practice infrastructure can and should evolve over time, will improve data searches and access to useful and allowed data, and enables data owners to easily expand or restrict data access as needed.

To protect and preserve both data and infrastructure, we recommend that the NAIRR develop guidelines that express the best modern security practices and access control principles, grounded in open standards and continually updated to reflect both improving technical standards and emerging risks in privacy and cyber-security. We suggest that the NAIRR adopt a risk-based layered approach to security and access control, where minimal risk data and resources are made highly accessible and widely available (e.g., to secondary school students) with minimal complexity for the user. In other words, the NAIRR should develop and adopt procedures to protect data, intellectual property and computing resources that are proportionate to the value and risk associated with those resources. A deliberate and integrated *DaaS* approach would enable strong, consistent, and flexible data and security policies, practices, and implementations that can evolve as technologies and AI evolve.

A comprehensive security design strategy provides security assurance through identity management – the process of authenticating and authorizing security principals. Identity management services should be used to authenticate and grant permission to users, partners, customers, applications, services, and other entities. A successful design strategy will also endeavor to classify, protect, and monitor sensitive data assets using access control,

encryption, and logging in the NAIRR data repository system, and will place controls on data at rest and in transit.

The evolution of standards for APIs, WebServices, metadata and so on make it possible to reduce dependencies on features of the specific platform that the data resides on. As data from different sources and different countries are contributed, each of those data sets may or may not come with different use and security restrictions, such as the European Union GDPR requirements. Because it will be impossible to identify all possible restrictions and requirements from the outset, establishing core requirements up front and building an adaptable system that can evolve with requirements over time is critical.

In addition, the NAIRR should play a leadership role in setting security and privacy standards in both academia and industry and provide appropriate technology and educational material to help researchers meet these standards. The NAIRR also should develop and deliver other educational materials regarding the use of NAIRR infrastructure and data sets, working with academic institutions and educational organizations to develop curricular materials that ensure students and researchers at all levels down to secondary schools can understand, access, use and analyze data.

Regarding Goal G - We believe the NAIRR can and should play a critical role in evaluating and defining research data practices and AI use and implementation guidelines to ensure equity and preservation of individual privacy and civil rights and civil liberties. Just because we *can* do something with technology does not mean we *should* do it; appropriate guidelines and guardrails must be in place to avert a “tragedy of the commons” when it comes to data and AI.

Accordingly, in addition to developing policies and practices for the work the NAIRR sponsors, we recommend that the Task Force investigate the broader social implications of AI research and technology beyond that work. We expect that the Task Force’s work will encompass ensuring equitable access to compute and data resources for researchers and educators at a wide range of institutions while also addressing the privacy and civil liberties implications of the use of data and the research enabled by that data. The Task Force should consider issues such as “rights to be forgotten” and challenges raised by differences in international law concerning data and AI systems and access to data and systems by researchers within the country of origin and researchers outside the country of origin.

Specifically, we recommend that the NAIRR should (1) form a taskforce that includes academic and government researchers, computer scientists, industry representatives, non- profit organizations, legal and ethics scholars who also have an understanding of modern computing technologies to identify the privacy, security, civil rights and civil liberties risks associated with the aggregation of massive amounts of data and the application and advancement of AI; (2) lead and engage public discussion of the risks created by data and AI;

(3) form specialized taskforces comprised of academic and government researchers, industry representatives and computer scientists to implement the recommendations of the first taskforce, including developing technical and organizational standards for privacy and

security in AI systems, potentially including standards akin to Human Subjects protocols for appropriately high-risk situations, and (4) develop and publish educational tools and guidelines for broad public understanding of the benefits of and risks inherent to data and

advancing AI.

The work to evaluate and maintain appropriate guidelines and policies will be an ongoing effort because changes in technology and the legal frameworks associated with data and AI will require regular appraisal of data sets and the AI systems that exploit them. While we do not believe the NAIRR or other branches of the government should be in the business of defining “representative data sets,” we do believe the NAIRR has a critical role to play in ensuring that data sets are appropriately and transparently curated and identified by metadata to enable researchers to evaluate the appropriateness of data sets for their specific projects and goals. Similarly, the NAIRR is well situated to devise standards for the implementation of AI to ensure responsibility, transparency, and that AI innovation is serving the public good and consistent with American goals, values and laws.

Regarding Goal H - given the importance of the NAIRR to maintaining the long-term global leadership of the United States in research, the long-term sustainability of the NAIRR as a public resource is essential. The increasing centrality of computational methods in research across disciplines and the potential solution to compute access that NAIRR represents will naturally attract public and private partnerships and funding. However, private entities may represent competing interests, and partnerships necessarily introduce complexity. Given that the NAIRR is intended as a public research resource, the foundational, sustaining funding should come from public sources. Private-public partnerships (PPPs) and private investment, both in cash and in kind, can provide solid supplemental funding aligned to specific and project-driven interests.

Because compute can be expensive to provide, and computationally-intensive research interests must be balanced with environmental costs of delivering compute, we encourage the Taskforce to consider tiered and subscription-based models for NAIRR use, along with rewarding researcher contributions to open resources (e.g., algorithms and research data sets) and providing technical assistance to optimize projects to use the NAIRR facilities most effectively. For example, we have found that researchers and start-ups need “white glove” infrastructure support to train large “Foundation” models on the Oracle Cloud.

In terms of sustainability, a public private partnership funding model, with the government providing the initial funding for the establishment and governance of the NAIRR, and the private and educational sectors contributing to the capital and operational costs, would ensure the long term viability of the NAIRR. Funding organizations such as the NSF and NIH could preference the use of NAIRR assets, creating a continual funding stream.

Regarding Goal I – the NAIRR should consider developing a heterogeneous infrastructure to support AI education and research from primary school levels through to advanced, university-based academic, government and commercial researchers. This involves provisioning diverse compute resources and data resources appropriate to the communities

that the NAIRR supports, as well as the infrastructure required to meet cybersecurity and privacy standards that the NAIRR requires. Rather than develop and maintain its own

infrastructure, we recommend that NAIRR utilize government, academic and commercial cloud computing and data resources from multiple vendors. This will enable the NAIRR to offer the latest computing resources to users. We also recommend that the NAIRR develop a common abstraction layer that enables users to develop AI systems in the same way across all vendors (“write once, run anywhere”). The NAIRR should adopt open standards for both data and compute resources to enable frictionless mobility from one vendor to another. As outlined above, this likely will require the NAIRR to take a leadership role in developing or extending such standards.

Developing an educational framework for AI is critical to both advancing AI innovation and ensuring responsible and ethical AI and data curation and use. The NAIRR should develop educational materials for primary and secondary students, undergraduate and graduate level students and researchers, and professional governmental, commercial and academic researchers. Importantly, these materials must be cross-disciplinary and include both the technical aspects of AI and the ethical, social and organizational issues involved in the development and deployment of AI. These materials can include a set of “best practices” studies demonstrating how AI models can be developed using NAIRR in an ethical, socially-aware fashion. The educational organizations of industry partners, for example, [Oracle Academy](#) and [Oracle University](#), as well as other private sector educational resources, and organizations like the Computer Science Teachers Association (CSTA), CS4All, AI4All, AAAI, ACM and IEEE can and should be important partners in developing educational resources. Of note, the ACM Education Council has been working for a number of years on [curricular guidelines at the intersection of computing, data science, and machine learning \(a form of AI\)](#).

Regarding questions 2 and 3, we believe the NAIRR presents an opportunity both to effect a significant increase in the quantity and quality of scientific research and to expand the number, diversity, and perspectives of researchers engaging in the active pursuit of new knowledge and discovery. Therefore, we propose a tripartite organization of the NAIRR to focus specifically on and support the following three distinct activities, to ensure continued US leadership in AI research and technologies, and address equity and diversity concerns in education about and access to AI technologies and related opportunities across all of society:

1. Provide computing and data resources to support academic research into novel AI technologies, novel applications of AI to other academic disciplines, and the effects and implications of AI and related activities, e.g., data collection, on society. This will involve providing very substantial resources at the scale required to train “Foundation” models from scratch, as well as the more modest resources required to apply such models to a diverse range of applications, including basic science, healthcare, finance, etc. Support for this basic and applied AI research is essential for maintaining international competitiveness across a wide range of industries.
2. Provide computing and data resources to support undergraduate educational activities at universities and 2- and 4-year colleges, integrated with appropriate educational materials

dealing with AI and related technologies, e.g., cybersecurity, privacy, etc., and the social and ethical implications of these technologies. Modern AI models require specialized hardware

(e.g., modern GPU-enabled computers) not available at many undergraduate institutions, especially non-R1 universities and 2- and 4-year colleges. It is critical that we make education in AI technologies available to all students at all educational institutions, for equity and diversity reasons, as well as economic ones. US industry has a deep and growing need for graduates with expertise in all areas of AI; to meet this need and remain competitive internationally, we must support AI education at the undergraduate level.

3. Provide educational and teaching resources at the primary and secondary school levels, integrated into math and science curricula, for teaching and learning about AI and related technologies, including interactive learning platforms. There is a fundamental lack of understanding and trust in AI and related technologies throughout society today, and without a broader social understanding of AI technologies, we run the risk that large sections of society will not be able or willing to take full advantage of the benefits AI can offer. Appropriate primary and secondary school educational materials, integrated into science and/or math curricula together with related topics such as statistics, could help students and teachers understand AI technologies and their risks and benefits. An interactive learning platform, perhaps based on a web notebook, together with appropriate teacher professional development courses and teaching materials, will need to be a primary focus of this activity.

Regarding question 4, there are existing organizations and resources that can serve as initial building blocks for both the research and educational missions of the NAIRR.

The considerable compute resources required for state-of-the-art AI research are substantial. For example, training one “Foundation” model carries cloud computing costs of about \$100K. Federal research funders, like the NSF and the NIH, provide grants in the same order of magnitude as these costs, so similar grant proposals and funding mechanisms could be used to support academic AI research. While there are multiple models for providing computing and other resources to academic researchers, we believe NSF’s CloudBank is the closest extant model to what we recommend for the NAIRR. Cloud computing infrastructure of the kind required to build advanced AI systems already exists within US industry, and this kind of infrastructure (rather than traditional supercomputers) is ideally suited to developing the next generation of AI systems.

Access and authentication to the NAIRR resource for researchers, regardless of funding source, would be managed by the NAIRR.

The compute resources required for teaching purposes in higher education are much less, and could be managed by a block grant to each educational institution. The allocated resources would be managed by the institution, and student accounts and access to the NAIRR would be governed at the institutional level, in accordance with policies and processes established by the NAIRR.

At the primary and secondary education levels, the required resources need to be planned and allocated as an integrated part of the regular curriculum, made available via the federal Department of Education to the relevant state-level educational authorities, who then would manage allocation and access in accordance with policies and processes established by NAIRR.

Regarding **question 5**, public-private partnerships (PPP) will be essential to the long-term efficacy of the NAIRR. AI-related computational research workloads have become sufficiently large and

complex that no single organization – government or commercial – will be able to adequately and efficiently service all researchers. For example, through Oracle for Research we receive requests for AI-related research projects requiring up to 1M high performance compute hours per year. Servicing these workloads is important for the United States to remain at the leading edge of AI research and innovation, and PPPs are a viable mechanism for achieving this. PPPs can engage in open standards development and implementation, support efficient networking infrastructure development and availability, and ensure the efficient and timely availability of computational resources for research and education.

In a PPP, multiple partners can combine their resources, infrastructure, and responsibilities, allowing the NAIRR to allocate resources more efficiently and thereby use some of these resources to provide compute to researchers from across the academic spectrum. Furthermore, by combining resources, infrastructure, and skills from members of the public, private, and educational sector, PPPs benefit from an economy of scale, leveraging multiple technologies.

At the same time, new models for PPPs must be developed; these models must take into consideration commercial, academic and public interests and strive for simplicity and efficiency in contracting and funding models. While the aforementioned NSF-funded CloudBank project may provide a starting point, it also exposes myriad challenges with creating a simple user interface and pricing model while accommodating different university requirements and commercial considerations and various relationships between various government organizations and private entities. Other existing PPP models to consider are those that have been designed to handle, protect, and share medical or personal data. Additionally, changes to government funding practices are needed to ensure that grant funding and related university overhead fees do not combine to create perverse incentives for researchers to elect less efficient computational infrastructure options.

Regarding **question 6**, as a new initiative, the NAIRR will be competing with existing public and private AI projects. The NAIRR will need to differentiate itself by providing large amounts of high quality compute and data resources to researchers at all levels of the academic spectrum.

The NAIRR must also position itself to address the dire shortage of compute resources among the educational institutions that train developers in AI skills. By focusing on this need, the NAIRR will help address the significant lack of opportunities to work with high performance computing among diverse communities traditionally underserved by the educational system.

The NAIRR must also broaden its services beyond the traditional AI research community. New kinds of AI technology (such as Deep Learning “Foundation” models) are significantly easier to use than earlier AI, and the NAIRR should make these user-friendly models available to non-computer scientists in fields throughout the natural sciences, social sciences, and the humanities.

Democratizing access to AI research, development and education requires the NAIRR to address both technical and organizational challenges. The NAIRR will need to ensure equitable access and control at the data level, the network/infrastructure level, and the application layer. It will also

need to overcome limitations related to identity and access management (IAM), privileged access management (PAM), and performance and bandwidth, as well as the ability to support the required diversity of access restrictions based on data and resource policies.

To balance the need to ensure wide and equal access with the need to protect privacy and data, Oracle recommends that the NAIRR adopt a “zero trust” approach to security and access control. Zero trust is an IT security approach that keeps sensitive data safe while staying compliant with privacy regulations. As the use of cloud services rapidly expands, there are new risks associated with compromised or stolen credentials of a privileged administrator or application and increased potential for data theft, and for cyber criminals to conduct cyber fraud, because effective security controls are often an afterthought. Zero trust makes it possible for organizations to regulate access to systems, networks, and data without giving up control. Therefore, the number of organizations that are moving to a zero trust security model (meaning trusting nothing) is growing, so that they can safeguard data with security controls that restrict access to the data according to a specific policy.

About Oracle

Our responses to the specific points of the RFI are informed by ongoing Oracle projects, some of which are described here to provide further context.

Oracle’s Implementation of AI across the Enterprise

Oracle has been using AI to solve hard industry problems for many years, infusing AI into some of the most widely used enterprise applications in the world. Oracle now makes its AI capabilities and platforms, including general-purpose deep learning “Foundation” models trained on the Oracle Cloud, available to developers and data scientists to use directly in their own applications. With this, Oracle’s AI Platform enables users to build truly intelligent systems rather than simply incorporating general purpose AI models, by embedding AI models in SaaS applications and by using the domain expertise, data and knowledge from user applications. Oracle’s AI Platform includes the following:

- Data Science Service – a managed cloud service that helps data science teams rapidly build and deploy machine learning models across every step of the data science lifecycle from data labeling to orchestrated pipelines for data science workflows
- AI Services – includes perceptual AI services for areas like computer vision, speech and natural language processing, and Decision Services for business scenarios like anomaly detection, forecasting and recommendations and built on top of Oracle’s Data Science Service to make it easy for developers to apply AI to their solutions without data science expertise
- Business AI Services – higher level, compound AI solutions to solve for specific business scenarios, such as the Oracle Digital Assistant conversational platforms or Document AI

Oracle Research and Development in AI

Machine learning (ML) is an important subset of AI, and the work of the [Machine Learning Research](#)

[Group](#) (MLRG) at Oracle spans a wide range of topics in ML and natural language processing (NLP), including aspects of ML devoted to fairness and privacy and how they apply to the business problems Oracle's product groups focus on; NLP systems for model-based sentiment analysis, named entity recognition, entity linking, co-reference resolution and product attribute extraction; and standard image recognition models and trying to understand how those models might help in training image recognition models for cervical cancer detection. In addition, the MLRG is responsible for one of the most scalable topic modeling algorithms that can train a model on billions of documents in hours on a cluster of computers, and is currently working on approaches to building large-scale, multilingual contextual embedding “Foundation” models that are resistant to many types of errors.

Also developed by [Oracle Labs](#), KeyBridge uses advanced machine learning techniques to detect complex, unknown, or zero-day activities of intruders in terabyte-scale, unstructured, or semi-structured data streams. Oracle is using [KeyBridge](#) to monitor activities against our most valuable assets in the cloud infrastructure, e.g., data exfiltration, reconnaissance, and brute-force identity attacks.

Oracle’s Engagement with the Research and Startup Communities

[Oracle for Research](#) works collaboratively with researchers worldwide to accelerate results and inform the development of tools and services for research. Researchers working with Oracle for Research have used AI and ML to accelerate drug discovery, support virtual reality molecular dynamics simulations, make advances in 3D pathology tools, and more. [Oracle for Startups](#) enables new companies to launch and grow on Oracle Cloud in a virtuous cycle of innovation. Oracle for Startups has worked with AI startups like [DeepZen](#), [GridMarkets](#), and [Aleph Alpha](#). The Eleuther research consortium and start-ups such as Aleph Alpha are using the Oracle Cloud to build open-source versions of GPT-3 style “Foundation” models, choosing to build their open-source “Foundation” models on the Oracle Cloud because it meets their extreme infrastructure needs and provides the best price-performance trade-offs.

Oracle appreciates the opportunity to provide input on this important topic. We know how to collaboratively develop and share infrastructure that creates opportunities for advancement and global leadership and are eager to work with you and others to create an accessible infrastructure that enhances and accelerates the position of the United States as a world leader in AI, across industries and research disciplines. We look forward to working with the Task Force, OSTP, NSF and other agencies and your collaborators to develop a national AI infrastructure that provides broad access, education and expanding opportunities for research and innovation in AI and beyond.

Federal Register Notice 86 FR 46278,
<https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Palantir Technologies, Inc.

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

EXECUTIVE SUMMARY

Palantir Technologies Inc. (“Palantir”) is a software company that provides data integration, analysis, and decision-making platforms. Our platforms are used as a data foundation and development infrastructure for artificial intelligence (AI), giving us unique insight into the challenges and successes of AI programs. Our software supports AI research on topics ranging from pharmaceutical drug combinations to COVID research.

We applaud the National Artificial Intelligence Research Resource (NAIRR) Task Force (“Task Force”) for identifying and advocating for a holistic computing ecosystem to support AI researchers, practitioners, and students via the implementation of a National Artificial Intelligence Research Resource. In our experience, data – the quantity, quality, and appropriateness to the problem – is the greatest determinant of an AI effort’s outcomes and should be central to the NAIRR’s ultimate implementation.

To that end, we encourage the NAIRR to focus on ensuring the completeness, integrity, trustworthiness, security, and proper use of data used for AI research by prioritizing infrastructure, governance, and intentional resource allocation.

Infrastructure and Data Foundation. Trustworthy and secure AI models require a trustworthy and secure data foundation (as the saying goes, “junk in, junk out”). In considering the NAIRR implementation roadmap, we urge the Task Force to prioritize specific infrastructure components that are most essential and practical to ensuring a quality AI output: intuitive data integration and pipeline transformations, audit logs, granular access controls, the ability to version and branch data, and collaboration features for annotating datasets and identifying addressable issues over time (e.g., statistical and other forms of unwanted data bias). With these key investments, the NAIRR effort has the potential to support, enrich, and grow the AI research community by providing a shared data environment, not just static datasets. The NAIRR data resource can enable tangible, trustworthy results through collaboration tools that allow for secure modifications, annotations, and improvements to datasets and the ability to share knowledge and performance metrics. These investments will enable the NAIRR infrastructure to scale and support cross-organization, cross-discipline research to springboard the U.S.’ AI capabilities.

Technical, Governance, and Cultural Awareness in Responsible AI. Not only should the NAIRR incorporate best practices around AI Privacy and Civil Liberties (PCL), Bias, and Ethics, but it can — and should — enforce these best practices (e.g., stringent security and access controls) through a combination of a) technical infrastructure that facilitates enforcement of key data protection and responsible AI principles; b) governance that incentivizes and rewards best practices; and c) cultural awareness and discipline-specific frameworks to contextualize AI research and guide determinations of when and how AI applications can best align with the interests of impacted communities. Through our decades-long work developing and implementing data integration and management platforms built on granular security and privacy-preserving capabilities, we have demonstrated that well-engineered data infrastructure must provide the technical implementation measures that attach to and reinforce critical institutional governance measures. The challenges that AI research will ultimately address are not exclusively

technological in nature, rather, they are techno-social and require an understanding of the cultural contexts, innovative data science practices, and institutional controls.

Problem Prioritization through Resource Allocation. The Task Force is well-positioned to provide oversight driving NAIRR use cases in an opinionated way towards pressing, relevant problems and historically overlooked or underfunded initiatives. Similar to how the National Science Foundation approves grants, the Task Force can provide incentives and allocate resources (compute, data, etc.) based on problem evaluation. For example, the Task Force could identify AI safety research as a topic requiring more focused research and development. When industry alone researches this topic, it is likely to direct its attention to narrow, commercially-focused use cases. The NAIRR, in contrast, could expand the scope and insights produced by the AI research community. With the appropriate infrastructure and governance in place, the NAIRR has the potential to shed light on how AI resources are being used, discover whether the most important problems (as seen through a broader societal impact lens) are being researched, drive investment, and compound knowledge from across academic, government, and industry participants.

Categorizing AI Research & Development

The themes above can manifest differently based on the category of AI research, development, and application in question. We recommend that the Task Force clarify which of these categories NAIRR aims to address, recognizing each has different data, infrastructure, and organizational needs.

- 1) ***Pure or Basic AI research and development*** focuses on advancing the state of the art of techniques and methodologies for artificial intelligence. Basic research requires access to compute, ability to build new network types, and scaling infrastructure for large datasets.
- 2) ***Applied AI research*** focuses on taking an existing AI algorithmic approach and exploring its utility in the context of a real-world problem. Applied research requires the ability to collaborate with non-coding subject matter experts (biologists, physicists, etc.) as well as access standard libraries (pytorch, tensorflow).
- 3) ***Operational AI*** productionizes use of AI using real-world data, with real-world outcomes. Operational AI requires DevOps tooling to deploy, scale, and monitor models.

OUR PROPOSED VISION FOR A NAIRR DATA INFRASTRUCTURE

We encourage the Task Force to commit to investing in the three areas mentioned above (infrastructure, governance, and problem prioritization) as part of the NAIRR implementation roadmap to give researchers and students access to a powerful, trustworthy, and secure end-to-end AI research environment. Elements of this vision are already a reality at NIH and other agencies (see our response to Question 5 on page 9), and the NAIRR should evaluate the feasibility and sustainability of incorporating these practices.

Ability to search for, select, and combine data: As stated in the introduction, data and the associated data infrastructure are the greatest determinants of an AI program's success or failure. The NAIRR should empower AI researchers to:

- **Easily discover potential training datasets relevant to their use case** through a transparent data catalog, as well as metrics and metadata on those datasets.
- **Branch a data set**, modify it, and make those modifications available to the broader research

community in a fully secure and transparent way. Branching datasets supports discovering, recording, and mitigating bias and other data issues for the entire research community.

- **Collaborate** on training data sets without compromising data lineage, integrity, or security.
- **Granularly secure data** with low-friction, built-in access control tooling. Not all datasets should be fully accessible to all users (e.g., to protect sensitive information) and oversight is required to assess the potential consequences of combining data sets.

Access to the platform through a straightforward registration process. New users should ideally be able to access and begin using tools the day that they sign up. A complex and/or drawn-out registration process will likely alienate some of the communities that the NAIRR is most seeking to serve. Automated security validations that are built into the technical infrastructure and a user-friendly interface will be key to enabling this step.

Ability to train, test, and retrain AI models. The NAIRR should provide researchers with the development environment, computational power, and tools required to train AI models in a streamlined manner, or enable them to use their own tools, connected via open APIs. Users should be able to easily capture and share their work products (libraries, models, data modifications) back to the NAIRR, enabling knowledge to compound over time. A technical infrastructure designed to capture and share knowledge enables appropriate contextualization of existing data and research, which in turn enables directing resources to priority problem areas.

Ability to evaluate AI model performance within the NAIRR platform. Researchers should be able to view and capture metrics that show the performance of their AI models as well as how a model compares to baseline standards, and to capture performance against different evaluation datasets. Researchers could have the option to make this performance data discoverable so that others can learn from their work, while also securing their data so that it is only accessible by those with appropriate permissions.

Ability to contribute back to the NAIRR. User friendly tools should be provided that incentivize researchers to contribute back to the broader NAIRR community. For example, researchers could make available new and improved data pipelines, data annotations, model templates, analysis frameworks, etc. Incentives such as additional compute allocations that encourage contributing knowledge back to the AI research community would allow the collective knowledge of the NAIRR compound over time and streamline future research efforts.

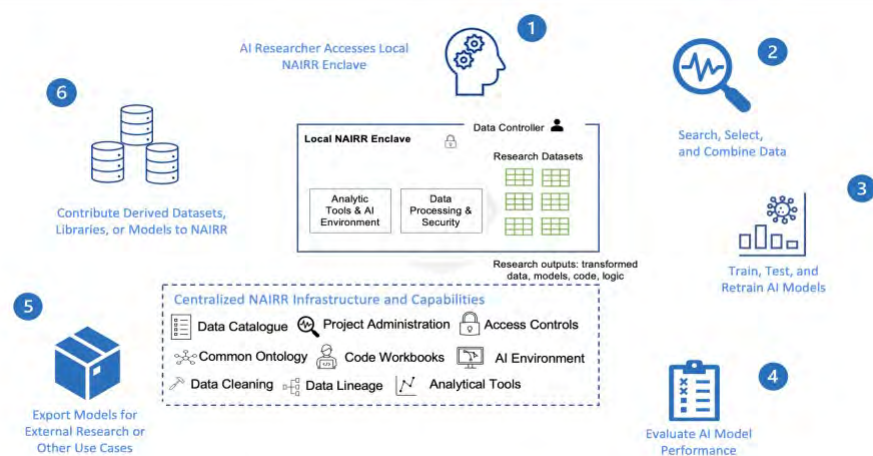


Figure 1: This graphic depicts the AI research lifecycle as enabled by a NAIRR digital infrastructure.

RESPONSES TO QUESTIONS

1.A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success.

In general, the “north star” for the NAIRR should be implementation of and provision of access to a data infrastructure that facilitates and democratizes AI research. Reflecting on our experience in enterprise data management, Palantir suggests the government consider the following NAIRR implementation goals:

Goal	Benchmarks
To provide the U.S. research community with high-quality, broadly available data sets that have been screened for common forms of statistical bias for AI model training and evaluation.	<ul style="list-style-type: none"> • Researchers have access to quantitative tools and techniques for assessing the quality of datasets and their effect on model outputs. • The NAIRR continuously integrates advancements in data evaluation tools and techniques.
To accelerate research of AI techniques and models that address critically important problem sets in both the public and private sectors.	<ul style="list-style-type: none"> • NAIRR infrastructure enables collection, reporting, and surfacing of metrics showing where research is being conducted. • NAIRR program office has access to these metrics and can ID research gaps.
To measure, and demonstrably improve, the effectiveness, accuracy, and fairness of AI models over time.	<ul style="list-style-type: none"> • The NAIRR platform contains infrastructure to calculate and evaluate the performance of AI models. • Metrics for AI effectiveness and accuracy are transparently and accountably recorded to curate the specific research area, along with other critical, contextual, and social relevant evaluations (e.g., fairness metrics)
To expand the AI research community to include researchers from historically disadvantaged communities.	<ul style="list-style-type: none"> • Metrics about who is using the NAIRR and what kind of research is being conducted are collected in the NAIRR platform (within the bounds of a consent framework and protected by access controls). Over time, metrics show increased participation from historically underserved communities.

1.D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country.

To create and maintain a shared AI research computing infrastructure that delivers access to quality and trustworthy data, we recommend that the NAIRR contain the following capabilities:

A dynamic, flexible, and secure data environment. Without integrated, clean data, model outputs are likely to be incomplete, inaccurate, and to perpetuate system flaws (including forms of bias). The NAIRR must have the capability for users of varying technical backgrounds to integrate any type of data source and require including critical information on data quality, provenance, and data pipeline health to protect against faulty input data.

Privacy and security controls. The Task Force should strive for maximum granularity in the NAIRR access controls, down to the row and column levels when required, to ensure that data will not be misused while also enabling the greatest flexibility in collaboration. Additionally, transparent propagation of security controls to downstream datasets will allow researchers to appropriately discover data used to build pipelines as required to pursue a properly scoped AI project. The NAIRR digital infrastructure should also contain a robust audit and logging system to ensure policy and regulatory compliance.

Collaboration functions, to include branching and version control. To allow for experimentation without requiring a separate development environment, the NAIRR platform should allow researchers to create, manage, and merge branches of their datasets, data pipelines, and models. This will not only accelerate and simplify the research and experimentation process, but allow for increased accountability and traceability of models, datasets, and data transformations within the NAIRR platform (e.g., ability to see when and where a new branch was created). Models and datasets could be directly compared in the platform to evaluate new techniques and identify sources of inaccuracy and bias. If a significant improvement in model accuracy, model fairness, or other metrics is validated within a branch, then researchers can share that back with the AI community, and leave annotations explaining the new branch. As researchers contribute to the shared data environment, the NAIRR will become the high-quality, validated, and collaborative data environment for AI model training and development.

User-friendly environment. To ensure the broadest possible participation within the research community, the user environment in the NAIRR should allow for easy manipulation of data even by non-technical users, and should include user-friendly security, compliance, and audit features for platform administrators. Given the existing technical barriers to entry within the AI research community, ensuring that the NAIRR is an inviting platform is critical to achieving the Task Force's goal of expanding the AI research community and democratizing access to AI tools.

Openness and interoperability. To achieve its full potential, the NAIRR must be compute-, storage-, and data format-agnostic. This ensures maximum flexibility for the government and broad participation from the research community. In compliance with data security and access controls, NAIRR users should be able to export their data, the code and logic in the data processing pipeline, the code responsible for building and running AI models, and the analysis code. This capability is critical to operationally deploying AI models outside the NAIRR or continuing research on external platforms using the same datasets and models.

1.E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets.

Based on our experience, successful data access and dissemination are linked with data infrastructure, including security and access controls, and stakeholder trust in the data infrastructure. We suggest that the NAIRR address barriers to sharing of high-quality datasets by avoiding a one-size-fits-all, all-or-nothing approach. Taking a more nuanced approach is likely to engender trust and confidence by government data holders and make them more amenable to sharing their data. While there likely are more than three relevant data categories, the examples below demonstrate the different levels of access control and data sharing policies the NAIRR should be able to accommodate:

- **Fully public data sources** with non-sensitive data such as ImageNet, Data.gov, and FBI Crime Data Explorer.

- **Enriched or annotated data sources** that may begin as public data sources or lightly-controlled data sources, but warrant policy and tooling to control their dissemination when enriched or annotated.
- **Combined data sources** formed by merging multiple controlled data sources or integrating researchers' external data sources with NAIRR-provided data. When combined, these data face many of the same issues as enriched or annotated data, and necessitate the option to evaluate them and control them appropriately.

The primary barriers to sharing more restricted, high-quality government datasets include:

Trust and Buy-In. Government agencies are less likely to share data if they cannot guarantee proper use and controlled dissemination. Technical infrastructure that enables secure data enclaves for multiple trusted partners to share the same data infrastructure can build trust and collaboration while ensuring security. See “Security” below and “Enable Multi-Stakeholder Engagement” on page 8.

Incompatible formats. The NAIRR platform should be interoperable to ensure that the widest possible selection of data can be ingested into the NAIRR platform for study. Additionally, the NAIRR should be able to automatically build and maintain live connections to relevant government source systems, so that models can be continuously tested on newly-generated data.

Security. To balance research security with inclusivity and flexibility, the NAIRR should include granular security and access control capabilities to allow for collaboration while ensuring compliance with privacy and civil liberty policies and regulations. A flexible framework for users to provide justifications to explain or qualify critical or risk-laden steps in their research workflows (e.g., data upload or export) should also be considered to reinforce propriety of user interactions with the NAIRR. Ensuring that the platform contains user-friendly and comprehensive administrator and management tools will also augment the ability for the research community to interact with government datasets.

Usability. Security control design can profoundly affect the usability and power of the platform for AI research. For example, Row-Based Access Controls, the ability for security markings and user permissions to be applied to each individual row of data, unlock many new opportunities for collaboration and platform accessibility. Instead of requesting access to an entire dataset or data pipeline, researchers may be able to use much of the information in a dataset while being restricted from certain elements (such as those that contain PII or other sensitive fields). The ability to automatically engage with specific subsets of data using granular security and access controls greatly expands the utility of the NAIRR and its accessibility to historically disadvantaged AI researchers. Additionally, the NAIRR should include a suite of deidentification tooling to further enable selective revelation of data, in compliance with necessity and proportionality considerations for data sharing, while mitigating risks of over-exposure and unintended privacy or data protection concerns.

1.F. An assessment of security requirements.

Security controls and requirements must be integrated into the foundational software fabric of the NAIRR, and should be automatically applied to not only data, but also to transformations, models, and combined datasets downstream of the original data. Please see our discussion of security in our responses to 1.E. and Question 3.

1.G. An assessment of privacy and civil rights and civil liberties requirements.

The NAIRR should adhere to relevant legal precedent as well as commonly accepted standards of privacy, civil rights, and civil liberties. This should include not only developing policy and governance standards for the platform, but also providing technical infrastructure inside the platform to help researchers maintain compliance with relevant standards. Cultural awareness considerations should factor into research efforts as well, both to help contextual the application of AI research to the complex social environments in which technology is actually ultimately used, but also as a means of anticipating ways that normative considerations may drive future legal and regulatory requirements. We provide further details on this topic in our response to Question 3.

2. Which NAIRR capabilities and services should be prioritized?

We recommend prioritizing a NAIRR data infrastructure that promotes and democratizes AI research. Please see our response to Question 1.D. above for tactical recommendations of priority capabilities to include in the NAIRR platform.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI?

Effective AI follows from the right data, properly managed. Through effective data integration and management, the NAIRR's policy and technical considerations should reinforce principles of responsible AI, which in turn will position the NAIRR to secure the trust of the AI research community and government data owners. To build trust, the NAIRR must be designed to provide visibility across data, processes, stakeholder organizations, and algorithms. In addition, NAIRR should provide mechanisms and toolkits to collect metadata on, understand, and remediate issues in underlying datasets that may contribute to forms of statistical or algorithmic bias. To reinforce and protect an ethical and responsible approach to AI R&D in NAIRR, Palantir suggests the following guiding principles:

Problem Prioritization, Review and Selection. There are certain problems that do not avail themselves of AI interventions. With the NAIRR providing a platform to democratize access to powerful tools and data, how the NAIRR evaluates and selects AI projects and datasets will be fundamental to the success of the program.

Methodically Assess and Address Sample Bias. The NAIRR can and should have the technical infrastructure to collect metrics about data quality so that it can assess and address sample and other forms of statistical bias. No matter how rigorous the design of an AI application is, if the data used for training and development is irredeemably flawed, the system output will almost certainly be compromised as well. This issue is even more pronounced where training data reflects systemic or other institutionalized forms of bias. As such, implementation of any AI system requires a clear assessment of the fidelity, quality, and representativeness of the data upon which its models are built and trained.

Ensure Auditability. To reinforce principles and practices of ethical AI, the NAIRR must provide accountability through automated traceability and auditing. In-platform system records and audit logs should be captured to allow for post hoc outcome analysis and documentation. The NAIRR is also in an ideal position to leverage the data collected in the platform to direct future research towards fundamental data or modeling gaps.

Ensure appropriate use with access controls. The platform should also include rigorous access controls that incorporate regulatory, legal, and normative expectations to provide the desired level of transparency, auditability, etc. Further details on our suggested approach to access controls are in our response to Question 1.D.

Enable Multi-Stakeholder Engagement. A critical operational principle for ethical AI usage is broad stakeholder involvement in building, deploying, and overseeing the NAIRR, as well as input in specific research selection and visibility into research outcomes. Multi-stakeholder engagement—particularly focused on communities impacted by AI—is critical to validating the NAIRR’s mission and building trust across the community of interested and affected parties. However, incentivizing multi-stakeholder engagement also requires trust and the guarantee that partners can maintain a level of security and control over their data to ensure proper use. Relating to the themes of data quality and transparency articulated elsewhere in our response, we suggest that the NAIRR’s technical infrastructure should be designed to enable multi-stakeholder engagement. This could include, but is not limited to:

- Infrastructure that supports secure enclaves within the same data environment.
- Collaboration and access control tools to allow the formation of data and research consortia.
- Visualizations attached to datasets to convey demographic composition intuitively.
- Built-in, automated compliance review workflows for sensitive data and research.
- Modeling results demonstrated with visualizations to promote understanding across various technical and non-technical stakeholders.

4. What building blocks already exist for the NAIRR?

While many potential building blocks exist around the government, such as ethical frameworks, the most important consideration for the NAIRR Task Force is to avoid duplicating efforts already taken by the public and private sectors. The government should leverage the billions of dollars of annual research and development currently undertaken by the private sector by engaging in public-private partnerships (see our response to Question 5 below for further details). For example, rather than building a new data infrastructure from scratch, NAIRR could partner with industry to use an off-the-shelf, continuously upgraded, open, cloud-agnostic software solution. With a technical infrastructure in place to allow for continuous delivery of new software features, the government would be able to not only maintain, but improve the NAIRR over time without major additional costs. We provide exemplars of private-public partnerships that can be used as models or building blocks in our response to Question 5 below.

5. What role should public-private partnerships play in the NAIRR?

Public-private partnerships are fundamental to ensuring the NAIRR becomes an effective shared research infrastructure, facilitating access for researchers and students to computational resources, high-quality data, educational tools, and user support. Recently, various U.S. Federal agencies leveraged innovative technologies to centralize and operationalize disparate COVID-19 data quickly, as well as provide platforms to enable a unified response to the pandemic. Similarly, the NAIRR will rely on public-private partnerships to empower a broader swath of researchers and students to perform cutting edge AI research. Additionally, the advances gained through NAIRR—while funded by Government—will provide benefits across the public and private landscape.

Some relevant capabilities and outcomes across public-private health partnerships include:

- Industry provides intuitive data integration and analysis software to lower technical barriers and promote broad engagement by making datasets reusable, referenceable, and auditable.
- Government and Industry partner to consistently implement [FAIR data principles](#), which could be similarly applied at the NAIRR, ensuring data is findable, accessible, interoperable, and reusable.
- Industry provides engineering data quality and security for Government research.

Palantir would also like to provide a few additional examples of our partnerships with the public sector to solve critical challenges through developing operational data assets.

U.S. National Institutes of Health (NIH) National COVID Cohort Collaborative (N3C) Consortium. Palantir software enables research across a variety of formerly disparate datasets and capabilities for the NIH N3C data enclave by automatically tracking data lineage of more than 5,000 data transformations across 65 sites, performing rigorous and automated data quality checks for trustworthy, research-ready data, harmonizing data at scale across more than 9.1 billion total records and over 3,000 users, and providing a collaborative workspace that has underpinned more than 30 scientific publications. Specific machine learning (ML) methods have been applied at N3C to improve understanding and response to COVID-19 and long COVID, including:

[Analyzing large datasets of clinical and demographic data to understand correlation between demographic characteristics and increased clinical severity of COVID](#)

[Leveraging ML to predict clinical severity and risk factors over time across a study of nearly 2 million patients and 34 medical centers nationwide](#)

U.S. Health and Human Services (HHS). Palantir software is the data infrastructure supporting the Department of Health and Human Service's Protect Platform. The Protect Platform assists the HHS and its partners to execute a holistic government response to fight the COVID-19 pandemic and protect the public's health, including providing a data infrastructure for analyzing public health data, such as [the effect of school mask mandates on pediatric COVID transmission](#).

UK National Health Service (NHS). At the outset of the global COVID-19 pandemic in March 2020, the NHS deployed Palantir software to help determine how to best distribute life-saving equipment, including PPE and ICU consumable items. The NHS uses Palantir software to bring together over 150 datasets from across the NHS and partner organizations to enable a unified data foundation and single source of truth. Sources include hospital supply chain, epidemiological, staffing, atmospheric, emergency call center, and COVID-19 test result data.

Beyond Palantir's direct experience, the NAIRR should look to other Federal programs striving to promote equitable access to AI tools and datasets. This includes NIH's nascent Artificial Intelligence/Machine Learning Consortium to Advance Health Equity and Researcher Diversity (AIM AHEAD) program.

CONCLUSION & KEY RECOMMENDATIONS

The Task Force can proactively craft processes and policies that will sustain the NAIRR over the long-term. Prioritizing data quality and policies that incentivize and promote flexibility and a constantly improving technical infrastructure will empower NAIRR with a sustainable program to harness the power of the U.S. technology sector to benefit the AI research community. We have summarized our key recommendations in the chart below:

Theme	Recommendation
Infrastructure & Data Foundation	<ul style="list-style-type: none"> ● Ensure researchers can easily discover training datasets, understand the data’s provenance, and transparently and securely modify or annotate datasets in a way that benefit other researchers. ● Provide tooling to allow researchers to contribute models, libraries, and updated datasets back to the NAIRR research community. ● Implement privacy- and security-enhancing controls directly into the NAIRR data infrastructure, including audit logging capabilities. ● Utilize open data, model formats, and documented APIs to allow researchers to export and import (while respecting permissions) to their preferred development environment. ● Emphasize collaborative capabilities within the NAIRR to include data discovery, data provenance and transparency, as well as branching and versioning of data sets.
Responsible AI through Technical, Governance, and Cultural Awareness	<ul style="list-style-type: none"> ● Adopt common standards around AI Privacy and Civil Liberties, Bias, and Ethics and enforce them through a combination of a) technical infrastructure to automate enforcement, b) incentivize and reward best practices, and c) cultural and discipline-specific frameworks to contextualize AI research and to guide determinations. ● Provide researchers with an environment and tools to assess model performance both quantitatively and holistically. ● Use NAIRR implementation to demonstrate effective U.S. global leadership for AI through ethical and responsible conduct. ● Drive trust in NAIRR through data security and access controls, carrying out periodic security audits, providing visibility across data and data set transformations, as well as ensuring transparency of active user groups and research efforts. ● Include diverse stakeholders (governmental, civil society organizations, researchers, industry, etc.) in periodic evaluations of NAIRR’s efficacy and responsible employment.
Problem Prioritization through Resource Allocation	<ul style="list-style-type: none"> ● Collect and analyze data about AI research focus areas; utilize incentives, including resource allocation, to shape those focus areas. ● Utilize adjacent governmental organizations, such as the National Science Foundation and the White House Office of Science and Technology, to help identify under-researched areas.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Partnership on AI

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



The Partnership on AI Response to the White House Office of Science and Technology Policy and National Science Foundation Request for Information: National Artificial Intelligence Research Resource

October 1, 2021

The Partnership on AI (PAI) is a non-profit partnership of academic, civil society, industry, and media organizations creating solutions so that AI advances positive outcomes for people and society. The Partnership on AI studies and formulates sociotechnical approaches to the responsible development of AI technologies to advance the public's understanding of AI and to serve as an open platform for discussion and engagement about AI and its influences on people and society. Today, PAI convenes nearly 100 partner organizations from around the world to be a uniting force for the responsible development and fielding of AI technologies. Partnership on AI staff composed this response based on some of PAI's recent work, much of it inspired and informed collectively by consulting over several years with our international group of multi-stakeholder Partner organizations. The information in this document is provided by PAI and is not intended to reflect the view of any particular Partner organization of PAI.

PAI develops tools, recommendations, and other resources by inviting diverse voices from across the AI community and beyond to share insights that can be synthesized into actionable guidance. We then work to promote adoption in practice, inform public policy, and advance public understanding. Through dialogue, research, and education, PAI is addressing some of the most important and difficult questions concerning the future of AI. Currently, PAI and its Partners work toward the responsible development of AI technologies in four Program areas: (1) AI & Media Integrity, (2) AI, Labor and the Economy, (3) Fairness, Transparency and Accountability, and (4) Safety Critical AI.

PAI is pleased to submit this response to the RFI for the National AI Research Resource on the specific questions referenced below. We will highlight PAI's work and publicly available resources on 1) Demographic Data and Algorithmic Bias 2) Responsible Publication Norms 3) Transparency through Documentation 4) Diversity, Equity, and Inclusion in AI and 5) Inclusion and Access to AI R&D through Multi-stakeholder Partnerships.

Question ID and IE:

“What options should the Task Force consider” on (ID) “including provision of curated data sets” and (IE) “[a]n assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource”

Demographic Data and Algorithmic Bias

The provision of curated data sets presents a challenge, particularly in collecting and using demographic data in service of detecting algorithmic bias with its many legal and ethical implications ([Demographic Data Convening](#), PAI January 2020). A lack of clarity as to the acceptable uses for demographic data is cited frequently by PAI Partners as a barrier to addressing algorithmic bias in practice. In order for data sets to be curated with an eye towards positively impacting historically disenfranchised groups, there must be awareness of what those groups are. This requires datasets to be disaggregated by race, gender, etc. in order to assess discrimination and inequality.

However, the inclusion of sensitive data is a fraught practice for datasets, especially with regard to categories that could be used for discrimination, such as sexuality, political affiliation, and immigration status ([Don't Overlook the Role of Demographic Data](#), PAI Blog, April 2020). Ensuring categories are representative of the populations in question requires an understanding of the constantly shifting nature of identity and ongoing engagement with marginalized populations. For example, the inclusion of the Asian Americans and Pacific Islanders (AAPI) category on US Census forms was a large bureaucratic struggle. While, at the time, it enabled representation of the AAPI community, over time many individuals categorized as AAPI have come to feel misrepresented. Disaggregated datasets should be managed by independent third parties, specifically ones that represent the at-risk groups reflected in the data ([Knowing the Risks: A Necessary Step to Using Demographic Data for Algorithmic Fairness](#), PAI Blog, September 2021).

The National AI Research Resource should encourage stakeholders developing and using demographic datasets to:

- Curate Datasets with support of community-based organizations that have trust and experience with the groups and communities in question.
- Ensure the need for demographic data to assess anti-discrimination, fairness, and inequality does not infringe on privacy rights or increase the undue surveillance of protected classes, vulnerable populations, or marginalized groups. (["What We Can't Measure, We Can't Understand": Challenges to Demographic Data Procurement in the Pursuit of Fairness](#), January 2021).
- Engage with the tensions around demographic data usage in AI and align with emerging work on the importance of equitable data as seen, for example, in the Executive Order On Advancing Racial Equality And Support for Underserved Communities Through The Federal Government ([Section 9](#)).

Question 3:

“How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?”

Responsible Publication Norms

A key development in reinforcing ethical and responsible research concerns establishing norms around the responsible publication of AI research. Over the past two years, PAI has conducted research and hosted multistakeholder convenings with the AI community to explore how advances in AI research can be disseminated in a responsible manner given their potential for misuse. Other research communities have established norms and procedures for publishing high-risk research such as bioethics and cybersecurity. However, the AI community has had less time to develop similar practices. The use or misuse of advancements in AI systems can lead to potential accidents, unintended consequences, inappropriate applications, and malicious uses but also contribute to potential systemic harms. AI can be used in biometric or facial recognition technology deployed in surveillance applications or bias in datasets and algorithms that are amplified when misapplied in different contexts like criminal justice or hiring.

In May 2021, PAI published a white paper titled, [Managing the Risks of AI Research: Six Recommendations for Responsible Publication](#), containing recommendations for the AI research community on responsible research and publication practices for anticipating downstream consequences. This report synthesizes insights gathered from our research and convenings, with over 30 individuals from PAI’s partner organizations and other stakeholders contributing to the paper. Shortly after its release, *Nature Machine Intelligence* [published an editorial](#) endorsing the white paper’s recommendations.

The key takeaway from PAI’s paper and surrounding engagement on this issue is that the AI community must build a responsible research culture that anticipates potential downstream consequences of AI research and mitigates risks. While our recommendations have primarily focused on interventions at the point of publication, there are other important, earlier stages in the research pipeline including conception, funding, access to computing resources/infrastructure, where actors can effect change to contribute to a responsible research culture.

The National AI Research Resource should consider the following three recommendations to reinforce responsible publication norms and foster a responsible research culture that includes a consideration of ethical and societal consequences by AI researchers.

- Disclose in publications

AI researchers must be encouraged to report the level of contribution their paper is making -- is it an incremental improvement or an entirely new technique -- and disclose the motivation behind the research so those evaluating can better understand its potential impact and develop mitigating strategies. AI researchers must also be encouraged to report the computation used in research projects to better understand considerations around reproducibility, and downstream consequences such as environmental impact.

- Normalize discussion about the downstream consequences of research

Researchers should be invited to reflect on the ethical and societal consequences of their research relative to the level of advances to the field provided by their work. For incremental advances, a short statement citing work that discusses the consequences of similar research in more detail may be sufficient. For more significant advances, a more substantial discussion is warranted. The goal of this exercise is to not only inform those evaluating the research including publication venues about potential negative impacts but to also encourage researchers to thoughtfully consider all the ways their research could be harmful, examine potential second-or third-order effects, and consider mitigation.

- Review potential downstream consequences earlier in the research pipeline

Research teams should be encouraged to build opportunities to consider societal impacts earlier during the research process including when initially formulating research ideas. In cases where the research could have high-stakes applications, or is likely to be a significant technical advancement, a more thorough review of societal impacts will be necessary. This exercise can take many different shapes including holding discussions with researchers, hiring cross-disciplinary experts, getting inputs from underrepresented inputs etc. While these processes lack the rigor of formal mechanisms like university Institutional Review Boards, they can still go a long way towards mitigating negative impacts and can be less burdensome for researchers involved.

Transparency Through Documentation

The Partnership on AI is working towards establishing new norms on transparency through documentation with our project on [Annotation and Benchmarking on Understanding and Transparency of Machine Learning Lifecycles](#) (ABOUT ML). The project identifies best practices for documenting and characterizing key components and phases throughout the ML system lifecycle from design to deployment, including annotations of data, algorithms, performance, and maintenance requirements ([How ABOUT ML Taps Collective Wisdom](#), PAI Blog September 2019). A goal of documentation for system deployment is to write down the socially salient

aspects of performance, including fairness, robustness, explicability, and other topics. Relevant and difficult-to-answer questions include what tests, monitoring, and evaluation have been done, and how does monitoring relate to social outcomes ([Operationalizing AI Ethics Through Documentation: ABOUT ML in 2021 and Beyond](#), PAI Blog, April 2021).

A core tenet of the project holds that documentation is important to consider as both an institutional process and an artifact because many teams and individuals have to incorporate completing and updating such an artifact into their work in order for it to be useful. ABOUT ML's goal is not only to recommend what information should go into documentation for all ML systems but also to recommend how organizations can effectively reshape their processes to enable the reliable completion and maintenance of documentation in an ongoing manner ([Section 1.1.2](#) ABOUT ML Reference Document).

The Partnership on AI researchers and Partners, along with other stakeholders and public commenters, have identified a need to create documentation for internal accountability. The need for accountability motivates organizations to invest in and build the internal processes and infrastructure to implement and scale the creation of documentation artifacts ([Section 2.2.1](#) ABOUT ML Reference Document).

A key component of this project is its Steering Committee, comprised of around 30 experts, researchers and practitioners recruited from a diverse set of PAI Partner organizations. The Steering Committee guides the process of updating ABOUT ML resources based on the public comments submitted and new developments in research and practice.

The convenings and conversations inspired by the ABOUT ML project have led experts in the space to converge on the need to modify the sets of questions and information to be shared externally. This need is based on the constraints of what organizations are willing to share and what information external stakeholders require to consider the ML system to be sufficiently transparent. There should be a broad and public conversation between organizations that build ML systems and key external stakeholders — including civil society organizations, policymakers, end users, and non-users impacted by ML systems — to determine what information would be necessary in documentation for external accountability.

Diversity, Equity, and Inclusion in AI

The lack of diversity in the field of AI has been well-documented. Simply put, diverse teams result in better outcomes that can issue-spot and control for many of the challenges AI researchers and developers face. As an industry, AI struggles to both recruit and retain team members from diverse backgrounds, particularly women and minoritized communities. Despite widespread awareness of AI's diversity gap, the crisis continues, and in spite of significant investments to address the issue, there remains a lack of clarity about which

initiatives work best. By investigating the pervasive challenges in ethnic, gender, and cultural diversity in the field of artificial intelligence, PAI's [Diversity, Equity, and Inclusion \(DEI\) Workstream](#) seeks to turn collected insights into actionable resources for those striving to make a more inclusive environment for people working in AI.

We view this work as paramount to the advancement of responsible AI. If we do not work to sufficiently address diversity and inclusion on the teams developing the technology, we risk compounding existing economic and social disparities experienced by women and minoritized individuals and communities. PAI has benefited from a [dedicated research fellowship](#) to advance diversity and inclusion in AI. The goal of this work is to learn both from the lived experiences of women and minoritized individuals in the field of AI and from those involved in DEI initiatives at organizations to share knowledge across the field about what the key challenges are and what solutions work.

The National AI Research Resource should consider the following resources to guide the development of transparency through documentation and encourage its community of stakeholders to take concrete steps and commitments to reinforce core ethical principles such as transparency and diversity, equity, and inclusion in the field of AI.

- [ABOUT ML Reference Document](#)

The goal of this Reference Document is to synthesize insights and recommendations from the existing body of literature to begin a public multistakeholder conversation about how to improve ML transparency. By providing a guide for practitioners to start taking transparency seriously, this document serves as a first step. A foundational resource, this living document includes an extensive literature review, suggested documentation sections for datasets, and surfaces current challenges of implementing documentation.

- [ABOUT ML Process Guide](#)

Because documentation is both a process and a set of artifacts, transparency and documentation need to be an explicit part of the discussion at each step of the workflow. The ABOUT ML Process Guide provides suggested documentation questions and considerations for each phase of the ML system lifecycle — from design and setup to observation and maintenance — compiled from the ABOUT ML Reference Document and academic literature.

- [Beyond the Pipeline: Addressing Attrition as a Barrier to Diversity in AI](#)

A [forthcoming study](#) conducted in-depth interviews with managers, people working in DEI, and workers who identified as belonging to historically excluded identities and analyzed themes from those interviews to get at the heart of the AI field's attrition

problem. The paper distills these learnings into a set of insights and recommendations that those working in AI organizations can take to improve upon their current DEI practices, beyond implicit bias and diversity training.

Questions 4-6:

“What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services? 5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model? 6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?”

Inclusion and Access to AI R&D through Multistakeholder Partnerships

Working with its multistakeholder community of partners in academia, civil society, industry and media, PAI is committed to increasing access, inclusion and participation in AI R&D.

One of PAI’s Partners, the Tech Policy Lab at the University of Washington, has extensive expertise in applying value-sensitive design approaches to technology policy. In 2019, PAI worked with the Tech Policy Lab to implement their [Diverse Voices methodology](#) within PAI’s ABOUT ML project. The aim was to solicit views and feedback from communities who are often the least likely to be consulted in the formation of machine learning system documentation practices that may impact them. The insights garnered through this consultation informed the inclusion of a glossary in the ABOUT ML resource library as well as the design and structure of the materials to promote clarity and navigational guidance to readers from diverse backgrounds.

Building from the lessons learned from the Diverse Voices team and the work of other responsible AI advocates, PAI launched the [Methods for Inclusion](#) research project that aims to enable AI researchers and developers to more effectively and ethically engage with a broad base of constituents and stakeholders in the development of their AI/ML projects. This work seeks to meaningfully include impacted communities in order to enable AI/ML developers to provide an array of products and services that can better meet the needs of diverse populations around the world, without further deepening existing social inequalities or generating harm. A forthcoming publication will identify a broad range of methodologies and practices that can be applied at different stages of the AI development process, drawing on the large body of scholarship that has grappled with the question of how to create inclusive channels of participation in other domains.

Drawing on these experiences at PAI, OSTP and NSF are encouraged to incorporate a focus on inclusion, participatory design, and democratizing access throughout all aspects of the development of the resource. Building a diverse community of stakeholders across sectors to

engage and inform the development and deployment of the resource over time will be central to its success.

Thank you for this opportunity to provide information about the Partnership on AI's work that could help guide the development of the National AI Research Resource. One of the benefits of multi-stakeholder organizations such as PAI is the opportunity to convene and connect diverse perspectives from across sectors, disciplines, geographies, and lived experiences - a critical component to understanding and developing a national resource. The Partnership on AI is happy to provide more details or additional information about the research, workshops, convenings, and other activities we conduct as we continue to develop resources and tools to prevent harms and promote the development of AI that benefits people and society. Please contact Rebecca Finlay, Acting Executive Director [REDACTED] and Mark Latonero, Senior Policy Advisor [REDACTED]

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Maria Patterson

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Response to RFI on an Implementation Plan for a National Artificial Intelligence Research Resource

Submitted by Maria T. Patterson, PhD

1. What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

B. If a truly democratized solution is established, an NAIRR should be completely reproducible at small scale by any entity.

C. Allocation of resources and the decision making authority is of high importance, especially if the NAIRR is a cross-disciplinary infrastructure.

H. Sustainability with partnerships with private sector like cloud service providers will be important but great care should be taken to prevent lock-in and reliance on any single provider or on any paid (and non open source) component.

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

- Usage of *only free and open source* technology and software, to democratize participation when researchers move work in and out of the system and to remain agnostic and prevent lock-in to private-sector services
- Compute resources in a self-contained environment that can be entirely reproduced when deployed on other infrastructure (e.g., using Docker images)
- Access to a shared registry service for snapshotting compute environments (with embedded data / models or persistent identifiers / pointers to external data / models) that is version controlled (e.g., Docker registry)
- A science friendly user interface that could also be used in the exact same manner on a local laptop (e.g., JupyterLab)
- Interoperability with other research infrastructure and ability to freely move data and resources from system to system
- Data search systems that can be architected by any researchers for domain-specific usage over shared datasets (i.e., multiple ways to search data for different purposes should be a capability)

- Persistent identifiers for datasets that are agnostic to physical data location so as to modularly separate the technical system for ease of upgrading
- Tiered levels of data storage - ephemeral and frequently purged for personal sandbox space, shared longer-term collaborative spaces, highly-curated and ID-ed, searchable datasets
- Capability for any researchers to publicly publish literate-programmatic “papers,” or fully reproducible peer-reviewed publications that allow others to recreate data analysis in its entirety. This could also include partnerships with research journals.
- Streaming data pipeline execution capabilities

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

For all human-centered research, regardless of whether or not researchers have collected publicly available data (e.g., increasingly popular social media mining, without explicit consent of individuals), an ethics review board (or “Institutional Review Board” IRB) should be utilized.

Equity, bias, fairness and transparency and accountability would be difficult to address in practice for AI without implementing a technical system for proactively monitoring models. While the adoption of “model cards” is one framework for documenting AI tools, this is a documentation solution that requires model providers to write this documentation, is only useful if other researchers are using the exact same model (limited applications in practice), and could still allow downstream users to adopt models or AI tools that are not appropriately used for their application. An overarching system that could continuously monitor “bias metrics” for evolving AI / ML models on different datasets would be a huge asset. An analogous service could be something like Kaggle submission boards with scoring and benchmarks.

Something similar to a “data donation” center where researchers could crowdsource voluntary contributions of donated data could allow for diversity of datasets and mitigate bias due to limited data on underrepresented populations.

Version controlling and timestamping models, data, publications, etc, similar to GitHub would be an asset.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

A model for an interoperable, shared technical infrastructure that co-locates data, storage, and compute resources with common analysis tools is a “data commons.” This has been successfully implemented for scientific researchers, collaborations, and data scientists and could be a framework for developing a national research resource focused on AI. (See “A Case for Data Commons: Towards Data Science as a Service,” Computing in Science and Engineering, Grossman, Heath, Murphy, Patterson, and Wells, 2016 at <https://doi.ieeecomputersociety.org/10.1109/MCSE.2016.92> or <https://arxiv.org/abs/1604.02608>, the Center for Translational Data Science <https://ctds.uchicago.edu/datacommons>, and Gen3 <https://gen3.org/>.)

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Private entities such as cloud service providers could and should provide a role for burst capacity when researchers hit their quota limits or need additional resources that cannot be provided by the NAIRR but should not be relied on for any core services that cannot be exactly substituted with open source software. The NOAA Big Data Project and its Cooperative Research and Development Agreement with large cloud providers is an interesting model that democratizes core access to data at no cost but allows private entities (cloud providers) the ability to charge for additional services. Looking to the open source community and successfully models their like managed enterprise services could be useful.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The “tragedy of the commons” risk that some researchers/collaborations/entities may dominate usage could be a problem. Quotas could be set within timeframes using a proposal system similar to a Time Allocation Committee process in the astronomical community <https://www.noao.edu/gateway/tac/>. Reviewing committees should be diverse and use best practices in peer review processes for removing bias from decision making.

Relying on private paid service providers (e.g., cloud service providers) for any component that would make it such that any researcher could not perform the exact same research on their own local machines (with appropriate resources) is an outright bad idea, both biased against small and underfunded institutions/organization and may also lead to vendor lock-in long term.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

**Savash Kapoor, Mihir Kshirsagar,
Arvind Narayanan**

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



September 1, 2021

Via email: NAIRR-responses@nitrd.gov

White House Office of Science and Technology Policy and National Science Foundation

Wendy Wigen, NCO,
2415 Eisenhower Avenue,
Alexandria, VA 22314

RFI Response: National AI Research Resource

Thank you for the opportunity to respond to the National AI Research Resource (NAIRR) RFI. We are academic researchers associated with the Center for Information Technology Policy (CITP) at Princeton University¹ and write in support of the Task Force's aim to create equitable access to the infrastructure that fuels AI research and development.

In our response, we highlight the significance of supporting a research infrastructure that is designed to independently test the validity of the claims of AI performance. In particular, we draw attention to the widespread phenomenon of the industry peddling what we call "AI snake oil" – promoting an AI solution that cannot work as promised.² Relatedly, we highlight how AI-based scientific research is often plagued by overly optimistic claims about its results and suffers from reproducibility failures. We submit that the Task Force's implementation roadmap for the NAIRR must include establishing a public infrastructure that can critically evaluate AI performance claims. This infrastructure is vital to the goals of the Task Force of ensuring that AI research serves our shared democratic values.

¹ In keeping with Princeton's tradition of service, CITP's Technology Policy Clinic provides nonpartisan research, analysis, and commentary to policy makers, industry participants, journalists, and the public. This response is a product of that Clinic and reflects the independent views of the undersigned scholars.

² *How to Recognize AI Snake Oil*, Arvind Narayanan, Nov. 18, 2019, available at <https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf>.

1. We need a research infrastructure that critically evaluates AI-based performance claims and ensures that those tools are designed to serve societal values. (*Response to Question 3.*)

Recently, the industry has converged on a troubling and widespread practice that applies the label of AI to applications that do not and cannot work. We dub this phenomenon of using a veneer of AI to lend credibility to pseudoscience as *AI snake oil*. The proliferation of AI snake oil in such applications is a distinct issue from concerns around bias, but is a major contributor to the negative consequences that result.

AI-based research has led to undeniable genuine and rapid progress in many domains, but it is important to distinguish between the classes of problems where AI tools have been shown to be effective. For example, AI has made significant progress in aiding with perception tasks, but it has struggled to predict outcomes involving complex social phenomena. Applications that claim to predict social outcomes but in fact do not have any predictive power are unfair even if they are technically unbiased, since they mask the fact that they do not work as promised and end up perpetuating outcomes that are not well calibrated to the needs. This is especially true when they are deployed in determining important life outcomes.

As an example, consider the AI tools that are purportedly designed to automate hiring decisions. The main claim made by many companies producing these tools is that AI can analyze body language and personality traits from short videos of candidates and function as “algorithmic pre-employment assessments” to make hiring decisions easier. While it is generally understood by experts that these tools cannot work and are usually no better than random number generators, that has not stopped companies from riding the AI hype and being widely funded and adopted. Raghavan et al. highlight that 18 companies working on algorithmic hiring systems have collectively raised over \$200 Million in funding over the last few years, though not all of these companies offer AI assessments of job candidates.³

³ Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. “Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices.” ACM Conference on Fairness, Accountability, and Transparency.

Similar claims prevail in a large number of applications where AI systems are claimed to predict social outcomes such as the likelihood of recidivism or identifying at-risk kids. But recent research shows that AI systems today are no better than simple rules at predicting social outcomes.⁴ However, this does not stop companies from marketing AI-based systems that claim to solve these problems, and as a result industrial applications of AI that purportedly predict social outcomes are proliferating. This phenomenon has a further pernicious effect of fueling the hunger for personal data for these fundamentally dubious applications of AI and giving rise to “black box” algorithms that cannot be explained. Furthermore, these applications tend to distract attention from designing more effective interventions.

As a result, we see evaluating validity as a core component of ethical and responsible AI research and development. The Task Force could support such efforts by setting standards for and making tools available to independent researchers to validate claims of effectiveness. The NAIRR could also help create oversight mechanisms and support efforts to regulate AI tools that are known to not work.

2. There is a reproducibility crisis in scientific research that relies on AI and machine learning that the Task Force should address. (Response to Question 3.)

Scientific research suffers from a closely related problem to the industry’s reliance on AI snake oil. Many studies that purport to rely on AI have results that are overly optimistic and lack reproducibility.⁵ But there are challenges in creating the incentives for researchers to independently and rigorously examine scientific claims that the NAIRR can help us overcome.

Evaluating academic claims about machine learning is challenging. First, the code tends to be complex and lacks standardization, which makes it difficult to understand and replicate models. Second, there are subtle pitfalls for researchers who fail to differentiate between explanatory and predictive modeling. Third, the hype and overoptimism about commercial AI often spills over into machine learning research and obscures the findings.⁶ All these, of course, are in addition to

⁴ Matthew J. Salganik et al. 2020. “Measuring the predictability of life outcomes with a scientific mass collaboration.” *Proceedings of the National Academy of Sciences* 117 (15).

⁵ Sayash Kapoor and Arvind Narayanan. 2021. “(Ir)reproducible Machine Learning: A Case Study.” Preprint available at reproducible.cs.princeton.edu.

⁶ Joelle Pineau et al. 2020. “Improving Reproducibility in Machine Learning Research (A Report from the NeurIPS 2019 Reproducibility Program).” arXiv preprint arXiv:2003.12206.

the pressures and publication biases present in all disciplines that have led to reproducibility crises.

Systematic reviews have started to identify reproducibility issues and overoptimistic results in many academic fields that are adopting machine learning methods (*see* Figure 1 below). But this is complex and expensive work. One estimate suggests that we spend over \$28 billion a year on preclinical research in the United States that is not reproducible.⁷ As machine learning methods spread across academic fields, focusing on the reproducibility of that research is critical to ensure its validity.

One of the major roadblocks to reproducibility research is that appropriate computing resources are difficult to secure. While researchers can rely on cloud services such as Amazon AWS, Google Cloud and Microsoft Azure for compute-intensive AI research, there are fewer resources available for those seeking to vet claims of performance. This problem has intensified with the shift of private firms undertaking research into new AI models. For example, natural language processing models routinely require large amounts of computational resources. But the cost of computational resources to replicate performance claims are often beyond the reach of independent researchers at research universities. This further makes reproducibility of research output by private companies inaccessible due to issues with data sharing and lack of access to computational infrastructure.

We recommend that the NAIRR prioritizes the support of systematic reviews of published research across fields adopting machine learning methods. For example, the NAIRR could establish and sustain a computational reproducibility infrastructure and serve as a reproducibility clearinghouse by setting up benchmark datasets for measuring progress.⁸ This would lead to significant strides towards the aim of promoting transparent, effective, and responsible research.

⁷ Leonard P. Freedman, Iain M. Cockburn, Timothy S. Simcoe. 2015. "The Economics of Reproducibility in Preclinical Research." *PLoS Biology* 13(6).

⁸ Benjamin Haibe-Kains et al. 2020. "Transparency and reproducibility in artificial intelligence." *Nature* 586, E14–E16.

Field	Paper	Year	Num. papers reviewed	Num. papers w/pitfalls	Pitfalls
Neuroimaging	Whelan et al.	2014	—	4	Incorrect train-test split
Autism Diagnostics	Bone et al.	2015	2	2	Biased evaluation; data leakage
Bioinformatics	Blagus et al.	2015	—	6	Data leakage
Nutrition research	Ivanescu et al.	2016	—	4	Incorrect train-test split
Text Mining	Olorisade et al.	2017	30	—	Multiple pitfalls
Clinical epidemiology	Christodoulou et al.	2019	71	48	Biased evaluation; data leakage
Recommender Systems	Dacrema et al.	2019	26	25	Weak baselines, don't share code or data
Toxicology	Alves et al.	2019	1	1	Multiple pitfalls
Computer security	Arp et al.	2020	30	30	Multiple pitfalls
Health care	McDermott et al.	2021	511	—	Multiple pitfalls
Medicine	Vandewiele et al.	2021	24	21	Incorrect train-test split, data leakage
Radiology	Roberts et al.	2021	62	62	Multiple pitfalls

Figure 1 [from Kapoor and Narayanan]: a list of systematic reviews that highlight overoptimism and irreproducibility in applied machine learning research across academic fields.

3. The NAIRR can promote effective data stewardship models for using datasets. (Response to Question 2, Item D.)

The creation of datasets has been pivotal in the development of AI applications. But there is an underexplored dark side to supporting the broad release of datasets without mechanisms of oversight or accountability for how that information can be used. The resulting harms include privacy risks and representational harms. The NAIRR can play a pivotal role in mitigating these harms by establishing and supporting appropriate data stewardship models.

Consider the challenge of “runaway datasets” as an example of a problem that the NAIRR might address. In the last few years, many datasets have been retracted due to ethical concerns. But our research has documented how, even

after retraction, these datasets can remain widely available and are used across the industry and in research labs.⁹ This phenomenon has been dubbed the problem of “runaway datasets.” Of course, the ethical issues that caused the researchers to retract the original dataset persists in AI applications that continue to use these datasets after retraction. This highlights the necessity of dealing with ethical issues throughout the lifecycle of the dataset instead of addressing ethical issues only when the dataset is released.

In particular, the existing ethical oversight mechanisms within academia such as IRBs (Institutional Review Boards) are poorly suited to deal with runaway datasets. “Human subjects research” has a narrow definition in the context of IRBs and thus many of the datasets and associated research that have caused ethical concern in machine learning would not fall under the purview of IRBs. Further, IRBs do not consider downstream harms during their appraisal of research projects.¹⁰ This compounds issues with runaway datasets and exacerbates ethical concerns with the creation and use of datasets.

The NAIRR can address this gap by creating centralized data clearinghouses to regulate access to datasets. Such clearinghouses could include safeguards for monitoring ethical concerns through the lifecycle of the use of the datasets. The NAIRR could also create a framework for licensing datasets and machine learning models so researchers can control the intended and acceptable uses of their work. For example, we see significant confusion resulting from the use of unclear and non-standardized licenses in dataset releases. Finally, the NAIRR could establish mechanisms for exercising responsible data stewardship that can make decisions about the ethical uses of datasets at the time they are being created and while they are in use. While some research projects already follow such a procedure when releasing datasets, institutional support including providing funding towards data stewardship committees would help reduce the ethical risks of AI applications due to runaway datasets.¹¹

* * *

⁹ Kenny Peng, Arunesh Mathur, and Arvind Narayanan. 2021. “Mitigating dataset harms requires stewardship: Lessons from 1000 papers.” arXiv preprint arXiv:2108.02922.

¹⁰ Jacob Metcalf. 2017. “The study has been approved by the IRB’: Gayface AI, research hype and the pervasive data ethics gap.” Pervade Team.

¹¹ Ian Lundberg, Arvind Narayanan, Karen Levy, and Matthew J. Salganik. 2018. “Privacy, Ethics, and Data Access: A Case Study of the Fragile Families Challenge.” *Socius*, 5.

We commend the Task Force's careful attention to these issues and welcome the opportunity to discuss any questions.

Respectfully submitted,

Sayash Kapoor

Graduate Student, Department of Computer Science

Mihir Kshirsagar

Technology Policy Clinic Lead, Center for Information Technology Policy

Arvind Narayanan

Associate Professor of Computer Science

Contact:

Website: <https://citp.princeton.edu>



Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

**John T. Feddema, David J. Stracuzzi,
James R. Stewart**

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Response to Request for Information on an Implementation Plan for National Artificial Intelligence Research Resource

John T. Feddema, David J. Stracuzzi, James R. Stewart¹

Sandia National Laboratories²

Sept. 1, 2021

Introduction

This paper is in response to the request for information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource (NAIRR) that would provide AI researchers and students across scientific disciplines with access to computational resources, high-quality data, educational tools, and user support. Sandia National Laboratories is supportive of the NAIRR and would like to participate in the direction and use of this resource. Below is a discussion of the questions posed in the RFI. The specific question addressed in the RFI is indicated in brackets.

Goals and Metrics [Question 1.A]

With recent improvements in specialized computing hardware (graphical processor units (GPU's) and Tensor Processing Unit (TPU's)) and Machine Learning (ML) algorithms, Artificial Intelligence (AI) promises to improve U.S. competitiveness, quality of life, and national security. A new NAIRR would accelerate technological advances in AI and enable its use in a broad span of applications. We envision that the NAIRR will develop fundamental capabilities that U.S. industry as well as many U.S. government agencies can build upon to meet customer and public needs.

The goals of the NAIRR should include:

1. Improve U.S. competitiveness
2. Improve U.S. quality of life
3. Improve U.S. Government access to AI

¹ Contact Information: John T. Feddema, [REDACTED] David J. Stracuzzi, [REDACTED] James R. Stewart, [REDACTED]

² Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2021-10794 O

4. Dramatically increase the number of U.S. universities performing AI research including the underserved.

Metrics that should be measured to evaluate success of NAIRR include:

1. The number of industry, university, and government partners who become members of the consortium
2. Total number of government and consumer products enabled by the consortium
3. The number of technical advances in AI research as measured by publications, open source and copyrighted software, and patents.
4. Volume and diversity of data accessible through NAIRR
5. Capacity and usage of computing resources available through NAIRR

Ownership and Administration [Question 1.B.i]

Advances in AI will have broad impact across many government agencies, including the Department of Energy (DOE), Department of Defense (DoD), Defense Advanced Research Projects Agency (DARPA), Department of Transportation (DOT), and Department of Commerce (DOC). At the DOE Office of Science, AI will enable improved scientific discovery in the areas of light water reactors, carbon capture, combustion research, bioenergy, energy storage, power grid, and advanced batteries. At the National Nuclear Security Agency (NNSA) and the three NNSA National Laboratories, AI will assist in the design, production, and surveillance of weapons systems; aid in the detection of weapons of mass destruction (WMD) proliferation; and protect US assets against cyber threats.

Because of the broad impact of advances in AI, we believe a Public Private Partnership and/or Consortium including multiple government agencies, universities, and industry should be considered. This consortium should be similar to the DOE Public-Private Consortia³ or FDA's Scientific Public Private Partnerships and Consortia⁴. The consortia would enable joint foundational research while leaving participants free to build on the shared information and software to create proprietary outcomes of value to commercial, public, and national security applications.

Since multiple government agencies could benefit from advances in AI, joint ownership by National Science Foundation (NSF), DOE, DoD, DOT, DARPA, and other government is desirable. If a single lead government agency is necessary, NSF should be considered since it aligns well with their mission to promote the progress of science and provide global leadership in research

³ <https://www.energy.gov/technologytransitions/downloads/doe-public-private-consortia>

⁴ <https://www.fda.gov/drugs/science-and-research-drugs/scientific-public-private-partnerships-and-consortia>

and education although DOE also has a proven track-record of leading large public-private consortia and the DOE governance model has attributes worthy of emulation.

Governance [Questions 1.B.ii and 1.C]

The governance structure could leverage the existing public private partnership governance structure from DOE. DOE proactively manages their consortia to maximize their impact and benefits and ensure a successful transition to industry. The DOE consortium management plan defines leadership and decision-making structures, methods for communicating among members, and a sustainability model. Roles, responsibilities, authorities, and accountabilities of key personnel are defined. The management plan enables members to plan R&D activities, share progress and discoveries, and evolve a roadmap to address anticipated needs. [Question 1.B.ii]

The governance board should include a subset of members from NSF, DOE, DoD, universities, and industry. The governance structure would include subcommittees with the necessary expertise to facilitate decision making. This would include decisions around membership, data standards (e.g., formats and privacy), computing ecosystem, prioritization of computing access, etc. It would be expected that the subcommittees be managed through term limits with new members being elected by the NAIRR membership or nominated by the governance board. Further, the NAIRR governance structure would be independently reviewed on a regular basis (perhaps once every three years). [Question 1.C]

Infrastructure [Questions 1.D – 1.G]

A commercial cloud infrastructure would be an excellent choice for computing resources. Grants from government consortium members, e.g., NSF, DOE, DoD, DARPA, would sponsor research performed on the cloud. Industry would join the consortium by either paying a consortium fee or providing commercial cloud infrastructure.

Note that while there are several well-established codebases for machine learning that cover core capabilities such as classification, regression, and clustering (unsupervised), less common ML algorithms and a majority of AI algorithms are available only as research code and may be unsuitable for non-experts. The infrastructure should allow individual researchers to customize specific algorithms in the shared codebase (locally) to support a broader array of applications than possible with standard implementations.

A separate piece of needed infrastructure relates to training material. If the expectation is for practitioners – people not specifically trained in AI and ML – to apply these algorithms to datasets relevant to their primary research areas, then they will likely require training beyond code documentation to support their efforts. AI and ML algorithms are typically not plug-and-

play. Algorithms must be selected carefully, and data must be appropriately prepared. [Question 1.D]

As consortium members, DOE and its National Laboratories might provide unclassified datasets in the areas of light water reactors, carbon capture, combustion research, bioenergy, energy storage, power grid, and advanced batteries. In many cases, these could be broadly shared among consortium members.

With respect to the training requirement noted above, a key barrier to use of data and compute resources may be wasteful experimentation. Consider for example deep neural networks, which require very large amounts of computation for training. A trial-and-error approach to determining network structure and hyperparameters may require many attempts before success (if it ever comes), each of which uses a large share of available computation. Hundreds of researchers iterating over these large trial-and-error loops can bog down both computation and storage of even a world-class computing resource. Great care and planning will be required to minimize unnecessary experimentation and focus ML training attempts on those that are likely to succeed. [Question 1.E]

Basic research should be performed at an unclassified level. This research would provide foundational software and mathematics upon which industry and government agencies could build products within their own computing infrastructure. The NAIRR network will need to be secure so that only U.S. consortium members have access. [Question 1.F]

Privacy of data must be preserved within the NAIRR cluster and network. Commercial cloud computing resources already have privacy built into their systems, although it is a good idea to check certification periodically. Any experiments including person information must be approved by a human studies board. Anonymized data is preferable as it would be expected that most research benefits would be gained through the sharing of data within the NAIRR amongst researchers without having to worry about privacy. [Question 1.G]

Federal Funding and Partnerships with Private Sector [Question 1.H and 1.I]

An initial allocation for each participating government agency in the consortium should be allocated by Congress. Each government agency would then be responsible for funding industries and universities to perform the AI research that is relevant to their needs. To minimize duplication, representatives from each government agency would share their plans and discuss how to coordinate overlapping research agendas. [Question 1.H]

NSF should be responsible for funding universities including and emphasis on Historically Black Colleges and Universities and other Minority Serving Institutions. DOE, DoD, and DARPA would fund collaborations with industries and universities, also emphasizing inclusion and equity. [Question 1.I]

Capabilities and Services [Question 2]

Highest priority should be the creation of a cloud infrastructure with access to datasets that can be used to test AI algorithms. Second highest priority should be access to educational tools and services. Third highest priority should be standards for measuring and assessing the performance of the AI algorithms on particular applications.

Ethical and Responsible AI Research [Question 3]

NAIRR will need to identify those applications where we must be concerned with racial and gender equity, fairness, bias, civil rights, transparency, and accountability. For some scientific applications such as inertial fusion or combustion research, these may not be a primary concern; however, for many applications that include personal information, this will be a major concern. For those applications, restrictions on the use of AI algorithms will be required.

Building Blocks for NAIRR [Question 4]

There are several building blocks that already exist for the NAIRR in terms of government activities. Within the DOE, NNSA, and Sandia National Laboratories, foundational research in AI is already taking place with an emphasis on DOE and NNSA applications. DOE's Advanced Scientific Computing Research office has initiated projects related to Scientific Machine Learning⁵, and NNSA's Advanced Simulation and Computing (ASC) Program has a new initiative in Advanced Machine Learning.

Scientific Machine Learning (SciML) has the potential to transform science and energy research. DOE has significant investments in massive data from scientific user facilities, software for predictive models and algorithms, and high-performance computing platforms that will benefit from advances in machine learning and artificial intelligence. Six prioritized research directions are

1. Domain Aware SciML – Integrating human expertise and domain knowledge with scientific machine learning methods
2. Interpretable SciML – new exploration and visualization approaches to interpret complex machine learned models using domain knowledge as well as metrics to quantify model differences
3. Robust SciML – research to show that SciML methods are well-posed, stable, and robust
4. Data-Intensive SciML – developing improved methods for statistical learning in high-dimensional SciML systems with noisy and complex data, for identifying structure in

⁵ <https://www.osti.gov/servlets/purl/1478744>

complex high-dimensional data, and for efficient sampling in high-dimensional parametric and model spaces.

5. Machine Learning-Enhanced Modeling and Simulation – developing new methods to quantify trade-offs and optimally manage the interplay between traditional and machine learned models
6. Intelligent Automation and Decision Support – new mathematically and scientifically justified methods to guide data acquisition and assure data quality, improved SciML methods for multimodal data encountered in scientific applications, and new methods to optimally manage resources using in decision support

Additional building blocks are being developed within the NNSA ASC program. The ASC Advanced Machine Learning (AML) Initiative will improve simulation capabilities in weapons design, production, qualification and certification activities, as well as stockpile assessment through advanced data-driven analyses. This will increase NNSA's agility, enabling greater exploration of design spaces and improved predictive capabilities while potentially lowering the cost of physics simulations and data analytics. Similar investments to the DOE Office of Science prioritized research directions are being made in advancing physics-constrained machine learning, improving our ability to employ machine learning with sparse data, validating and explaining machine learning, exploring learning hardware in a high performance computing environment, creating AML-tailored data environment, and improving simulation workflows.

Finally, the Computing and Information Sciences Research Foundation within Sandia National Laboratories, a DOE multimission NNSA laboratory, has recently started an initiative in Trusted AI. This strategic initiative focuses on the rigorous foundation required to use AI technologies and advancements in high-consequence national security applications. Three thrust areas are being explored:

1. Mathematical Foundations – Improvement of AI methods, relying heavily on abstraction and theory. Emphasis is on understanding the power and limitations of AI methods, creating novel approaches to training and inference especially with limited data, and ensuring robustness especially in the context of extrapolation
2. Efficient and Secure Systems – AI tools, algorithms, methods, architectures, and hardware that mitigate open research challenges in scalability, data management, domain and architecture-awareness, and counter-adversarial security
3. Usability and Trust – Improved decision-making performance and understanding of trust and use in national security applications. Emphasis is on developing a causal, theory-based understanding of trust and use, including when and why they dissociate. This includes creating trust measures necessary for making decisions based on AI results, principled approaches to domain-informed AI that increase user trust and understanding adversarial impacts on overall decision quality.

All three of these efforts plus similar efforts at the other DOE National Laboratories will want to participate in the NAIRR consortium. Many of these efforts are already collaborating with U.S.

universities, and the new NAIRR consortium will allow these current efforts to leverage similar research efforts being sponsored by other U.S. government agencies

Potential Limitations [Question 6]

There are several potential limitations that could prevent the NAIRR from democratizing access to AI R&D. First, industry leaders (Google, Amazon, Facebook, Microsoft, IBM, etc.) may not want to participate in this consortium because they are already working with universities and may not find the consortium model to be a competitive advantage to them. Allowing these industry leaders to provide cloud computing resources as an in-kind membership fee will help draw them into the consortium.

Second, rules on the sharing of intellectual property could also be a bottleneck and delay initiation of consortium research activities. As found in Sandia's Combustion Research Facility⁶, a variety of consortia models will need to be established to meet the needs of the partnering organizations. Also, the consortium should be focused on advancing fundamental knowledge. Intellectual property and proprietary designs can be created by industry after the consortium's development of the fundamental knowledge.

Third, public funding could also be an issue. A path for sustained funding across government agencies will be a critical part of ensuring continued democratized access to AI R&D.

Fourth, as mentioned previously, AI algorithms are not plug-and-play, regardless of how easy any given software implementation is to use. Most successful AI applications entail extensive use of both domain expertise and AI expertise. The point of developing a national AI computing resource therefore needs to include matching AI expertise with domain expertise in the application spaces to ensure that AI algorithms are being applied appropriately, and that any available domain expertise is incorporated into attempted solutions. If use of a national AI computing resource devolves into an unmanaged set of AI applications that ignore either relevant domain knowledge or mathematical limitations of the algorithms, then the NAIRR will not produce the desired technical advances.

⁶ <https://www.energy.gov/sites/default/files/2015/09/f26/CRF%20Case%20Study%2008-11-15%20FINAL%20CR.pdf>

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

SAS

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



Executive Summary

Establishing a National Artificial Intelligence Research Resource (NAIRR) is an ambitious and admirable goal. With the NAIRR's envisioned computing and data infrastructure to democratize AI, collaboration between various stakeholders will lead to better research, innovation, and, ultimately, citizen services. SAS Institute Inc. (SAS) welcomes the chance to provide suggestions to the NAIRR Task Force to help shape its roadmap strategy. In our more than 45+ years as a leader in the field of advanced analytics, SAS strongly agrees with the goal of democratizing access to actionable AI.

It may be surprising that SAS' view is that NAIRR should embrace both a variety of open source and commercial tools and methods as a commercial software vendor. In our view, democratizing analytics requires data scientists to collaborate across coding languages (e.g., SAS, Python, R, Lua, etc.) as well as low-code analysis techniques. Successfully accomplishing this means NAIRR should pursue a strategy built upon choice rather than one around a preferred technology solution. The Task Force should consider deploying solutions that have capabilities to enable researchers to detect and mitigate bias (both in data and in models), explain and interpret the output of their models at a global and local level, and employ privacy preserving techniques.

Unleashing the power of an open analytics ecosystem must balance this choice of techniques with the need for a common framework for collaboration, governance, data management, and explainability. This layer of control will allow researchers with different skill sets to collaborate on a common problem seamlessly, ensure the data being run through AI algorithms is high quality, and ensure that models being run across different languages can be explained effectively. Governance of data and analytic models is also a key component to reinforce principles of data ethics and responsible research.

Blending this strategy together at scale requires interconnected initiatives related to people, process, and technology; something best done through public-private partnerships (PPP). SAS alone is investing \$1 billion in AI over three years through software innovation, education, expert services related to topics ranging from advanced analytics, machine learning, deep learning, NLP and computer vision. Leveraging these types of investments across the private sector will increase the speed to reaching a successful implementation of the NAIRR strategy. SAS is eager to make meaningful contributions to this effort and welcomes any follow-up to our recommendations.

SAS Feedback on NAIRR Roadmap

1. *What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]*

Roadmap Element D – Capabilities to create/maintain shared computing infrastructure
This element will require consideration of the people, processes and technology used to enable the NAIRR.

People considerations

To enable research, especially in traditionally underserved communities, institutions, and regions, the Task Force should consider using solutions that:

- Are well supported by training resources, online communities, and technical support.
- Have a significant user base so that developed skills are easily transferable.
- Provide graphical user interfaces in addition to coding interfaces – enabling domain experts, citizen data scientists, and users who may be new to the field to take advantage of AI capabilities faster, without having to become full-fledged programmers. In fact, experienced researchers may find it beneficial to be able to test ideas in a graphical interface without having to write code.
- Enable use and provide support for people with disabilities.

NAIRR leaders and administrators should have a firm understanding of data sources utilized, and possess and/or have access to those with domain expertise in AI, machine learning, data governance and infrastructure knowledge associated with deploying AI solutions (e.g., hardware, networking, cloud architecture, etc.).

Process considerations

Processes are critical to consider roadmap element D. The Task Force should:

- Develop processes for open-source package management.
- Consider how easy it will be to apply system updates for COTS and open-source solutions.
- Establish processes for ensuring that the NAIRR is restricted to use by the intended audience and appropriately shareable amongst those with access.
- Leverage automation wherever possible to promote repeatability and ensure all resources are being used efficiently.

Technology considerations

Finally, as technology is the core of NAIRR operations, the Task Force should consider:

- The variety of software and hardware used to develop AI applications. There are

many open source and COTS products that can be used to develop AI applications. Many times, researchers may want to use multiple programming languages and hardware as part of an overall development approach. Similarly, teams of researchers may have different skill sets on the same team, yet still need to collaborate on a common problem. Given this, the NAIRR should be a platform in which researchers can seamlessly apply a variety of languages and techniques using both open source and COTS products. This will have the additional benefit of not locking AI researchers into any specific technology.

- Scalability will be critical. AI projects typically ingest massive amounts of data and are computationally intensive. Computational constructs that leverage containers and Kubernetes can enable elastic scaling and can provide workload management features, so researchers are not competing for compute resources. These constructs can also provide fault tolerance in the event of a hardware failure.
- Access to specialized hardware and optimized software. Some types of AI algorithms, such as those used for deep learning, run more efficiently on specialized hardware chips like GPUs (as opposed to traditional CPUs), accelerating the time to train models. Many AI applications also require millions of complex transactions per second which benefit from optimized software that takes advantage of techniques such as in-memory computing. The Task Force should consider providing a platform that allows for access to the specialized hardware and software required for modern AI research.
- Centralized administration and governance. The Task Force should consider solutions that provide activity tracking of various aspects of the NAIRR environment, such as servers, job content, and usage. Aside from making the environment easier to maintain, this will facilitate a chargeback model if that is something the Task Force chooses to implement.
- Ecosystem integration. Given the wide variety of technologies that might be included in the NAIRR, the Task Force should consider how these will interoperate. Open software that enables communication by REST APIs will enable technologies to be more easily integrated.
- Security. The technology should have authentication and authorization mechanisms to ensure appropriate access to technical capabilities and data sources. All-or-none access permissions placed on datasets may protect sensitive information – including personally identifiable information (PII) and personal health information (PHI) – but unnecessarily restrict the utility of data available for AI applications. The Task Force should pursue a solution that includes automated detection and masking of sensitive information and/or row-level security permissions to ensure a larger, more representative amount of data is available for researchers. The Task Force may also want to consider whether to encrypt sensitive data and ensure its secure transmission.

- Data ethics capabilities. The Task Force should consider deploying solutions that have capabilities to enable researchers to detect and mitigate bias (both in data and in models), explain and interpret the output of their models at a global and local level, and employ privacy preserving techniques.
- Data curation. An AI application is only as good as the data that goes in it – and it provides the best return when it is supported by a well-governed data management program. AI systems do not merely extract insights from the data they are fed (as traditional analytics do) – they actually change the underlying algorithm based on what they see in the data. The more data they are fed, the more tightly they define the algorithm and the more confidently they make classifications or predictions. Because of this feedback loop, errors can multiply upon themselves if bad data and/or biased data (see Q3) is fed into the AI application. The dangers in that are obvious: inconsistency, inaccurate insights, loss of trust and decisions made that are misaligned with values and policy.
- Data management. In addition to providing curated data, the Task Force should consider data cataloging solutions to make it as easy as possible for the research community to find curated data, identify data quality issues and other attributes of the data that may need to be addressed, and evaluate its fit for purpose – including the potential to introduce or amplify bias in the resulting models.

2. *Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?*

In terms of prioritizing the items above, we would recommend ensuring people with the requisite skills are in place first. It would be valuable as well to coordinate and share learnings that emerge from bringing researchers together from academia, the private for-profit and nonprofit sectors, and the public to use those learnings to adapt to the technological and human infrastructures that advance multi-, inter-, and trans-disciplinary research. With these inputs, the Task Force would then be able to develop thoughtful processes and implement the required technology.

Once operational, NAIRR administrators should monitor usage to determine whether there are topic areas that are more addressed than others, commonly encountered problems, and other aspects of the environment for insights into subsequent prioritization of capabilities, services, and the curation of data sources.

3. *How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?*

AI and related technologies do not exist in a vacuum. They are inextricably linked to society and the natural world around us – often influencing one another in unexpected ways. Researchers leveraging the NAIRR platform should be provided with the necessary resources to practice developing AI technologies responsibly. The NAIRR task force

should consider a four-pronged approach to reinforce principles of data ethics and responsible research:

1. *High Quality Data / Data Governance*

Many instances of AI products and insights that have gone awry stem from the utilization of data that do not appropriately reflect the breadth of experiences within the larger population, or that perpetuate biases, or are not fit for purpose. The Task Force should consider the following:

- The data made available to researchers on the NAIRR platform should be carefully considered, managed, and governed appropriately with the goal of improving quality over time.
- Data should be representative of the diversity of the country as a whole, be granular enough for insights to be derived at a local level, and include a rich variety of high-quality data sources that could cover a broad range of research interests.
- Researchers should have access to sophisticated data matching techniques that would help them connect and enrich data sources.
- Data sets should be accompanied by datasheets that provide transparency to researchers regarding how the data was collected, its limitations, and any other ethical considerations.

2. *Technology*

The NAIRR platform should provide researchers access to analytical tools to reinforce principles of ethical and responsible innovation which may include, but are not limited to, the following:

- **Detect and Mitigate Bias.** Capabilities to detect and mitigate bias in data and in models are becoming more widely available. These include, but are not limited to, the ability to assess appropriate representation in data, flag sensitive features, identify proxies for sensitive features, assess model performance differences by sensitive features and their proxies, assess model performance differences within the feature space (independent of sensitive features), offer a range of fairness definitions, and bias mitigation algorithms for selected fairness definitions.
- **Explain and Interpret.** Capabilities to understand how complex models behave overall (globally) and at the individual observation level (locally) could include, but are not limited to, surrogate model interpretability, explainable machine learning models, natural language explanation of model results, and causal inference.
- **Consider Privacy and Security.** Capabilities to empower researchers to respect the privacy of data subjects should be offered. These may include the ability to automatically flag protected data and apply privacy preserving techniques to the data

(e.g., differential privacy, encryption, synthetic data generation, etc.).

3. *Ethics training for researchers*

Technology alone cannot solve for equity and other ethical considerations. Researchers should not only ask whether they *can* build data driven systems and insights, but rather whether they *should* in the first place. It requires thoughtful consideration of the impact their work can have if replicated at scale. Familiarizing researchers with data ethics principles of human-centricity, transparency, inclusivity, accountability, robustness, and privacy will be a critical first step to empower them to acknowledge blind spots and build data-driven insights responsibly. Mandatory training should be provided to researchers that emphasize these points before leveraging this platform. Training should also focus on how to translate principles into practice by providing tangible best practices on how to interpret results, detect and mitigate potential biases, and consider privacy and accountability throughout the entire data and model life cycle.

4. *Network of diverse perspectives*

The onus of developing AI technologies and insights responsibly and ethically should not solely fall on the shoulders of researchers leveraging the NAIRR platform. Researchers should have access to a wide and diverse set of perspectives in an effort to mitigate their own blind spots and biases. This will be especially important for researchers who intend to work on high-risk use cases (i.e., researchers could be required to submit a proposal that covers what they intend to do within the platform and a method of determining risk would need to be developed). The NAIRR Task Force should consider providing access to multi-disciplinary and diverse voices to researchers using the platform. Offering researchers access to diverse stakeholders can be enabled through public-private partnerships, partnerships with minority serving institutions (e.g., HBCUs), non-profit entities focused on advancing data ethics and responsible AI, and organizations that aim to bring the voices of impacted communities to the table. The following areas of expertise are recommended:

- Subject matter experts of the data being provided on the platform
- Social scientists
- Data ethics practitioners
- Ethicists
- Equity researchers at the intersection of various domains (e.g., health equity, transportation equity, education equity, etc.)
- The voice of impacted community and stakeholder groups

Researchers should be encouraged to consult with the network of diverse stakeholders as they formulate their research topics, identify candidate data sets for analysis, clean and analyze data, build models, and derive insights from their analysis.

4. *What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?*

Given one of NAIRR's goals is to enable all of America's diverse AI researchers to participate in advancing AI, the Task Force should consider partnerships with minority serving institutions like Historically Black Colleges and Universities (HBCUs) that have developed data science and machine learning curricula at both the undergraduate and graduate levels. Providing these students and researchers with access to rich data sources and NAIRR's robust AI R&D infrastructure will be a critical ingredient to ensure the AI workforce reflects the diversity of America.

There are several examples of research platforms (powered in part by SAS) that facilitate the types activities, resources, and services desired for NAIRR. Two include:

- [The National Opinion Research Center \(NORC\)](#) (out of the University of Chicago) manages a data enclave that is a high-performance computing environment with cutting-edge statistical, analytical, visualization, data management, and reporting tools. Since 2006, state and federal agencies, research institutions, foundations, and universities have used the enclave to securely house and provide remote access to confidential data. Enclave-based research informs a wide spectrum of public and private sector decision-making, as well as journal articles, books, position papers, conference presentations, dissertations, etc. At any given time, the enclave supports over 1,000 researchers via contracts and grants with a wide variety of government, academic, nonprofit, and commercial clients.
- [Project Data Sphere](#) is an independent initiative of the CEO Roundtable on Cancer. It leverages experts spanning industry, academia, and government to achieve mutual goals of improving cancer trials in order to expedite drug discovery. As the leading oncology open access data sharing platform, Project Data Sphere hosts de-identified patient-level data contributed by industry, academia, and PDS research programs. By openly sharing data, convening world class experts, and collaborating across industry and regulators to catalyze new scientific insights, Project Data Sphere accelerates delivery of effective treatments to patients. An open-access data-sharing model gives researchers rapid and ready access to data sets and analytical tools.

5. *What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?*

Public-private partnerships (PPPs) bring forward the benefits of public-sector priorities to coordinate collective action together with the innovation capacities and efficiencies of the private sector, including both for-profit and nonprofit organizations. Within the Artificial Intelligence space, the National Science Foundation-funded National Artificial Intelligence Research Institutes demonstrate the abilities and willingness of publicly funded agencies to work in partnership with private sector organizations to jointly fund research to be done by consortia of public and private institutions of higher education and nonprofit research and development organizations. The Institute of Education

Sciences has funded state agencies to work with private sector service providers to develop statewide longitudinal data systems, some of which SAS has worked with and incorporated AI- and machine language-powered data management and analytics. Additionally, SAS served to advise states on developing new governance, including ownership, structures, and processes, that oversee and operationalize new public-public partnerships between state agencies. There are just a few of the countless examples of federal agencies directing collective interests and action through grants, contracts, and cooperative agreements.

PPPs have the potential to create new capabilities and capacities, particularly through multi-, inter-, and trans-disciplinary collaboration to better identify and define outstanding questions and the lenses by which solutions are ideated and pursued. SAS' experiences in advising new partnerships, facilitating those partnerships, and being a PPP participant have taught us that effective PPPs have clearly defined common objectives, mutual benefits to participants, and shared resources (e.g., financial, human, physical). Our PPP experiences in education, energy, transportation, water, agriculture, defense and public safety, and health sectors were guided by clear principles that set the directions of the partnership in the form of enabling constraints that maintained flexibilities to better allow for emergent work, behaviors, and dynamics. Less successful PPPs have often been characterized as having highly prescriptive rules that served to overly constrain the activities of the partners and limited the potential innovation and novelty that comes from connecting expertise of the different partners. SAS has continued to accumulate the learnings from our guidance, facilitation, and participation in PPPs across several disciplines and industries, including learnings that serve to accelerate and amplify the benefits of PPPs while also dampening or avoiding those factors that lead to more unproductive pathways.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

According to a [McKinsey survey of 1,000 leading executives](#) managing analytics initiatives, only 8% had successfully scaled practices from the pilot to production stages. Although the NAIRR platform is intended for research and will not be a true "production" environment, the underlying infrastructure must be production-level quality to mitigate any risks to a truly democratized AI platform through successful management of the shared analytic resources (e.g., curated datasets, models) and scalability of the compute resources.

In SAS' extensive experience working with academia and government, the false dichotomy pitting open source analytics tools, such as Python and R, against commercial off the shelf (COTS) platforms, such as SAS, DataBricks, and IBM, creates unnecessary choices. The dichotomy drives the disconnect between the promise of democratized AI and successful, trusted results. Failed AI projects typically choose a strategy heavily

dependent on a specific set of tools and then face the consequences when the results are either disjointed and ungoverned or closed off and vendor dependent.

SAS has learned that successful projects embrace the dynamism, access and choice offered through open source tools while also gaining the governance and decreased time to value offered by commercial platforms. This combined industrialized analytics approach can offer democratized AI in a manner that benefits the field quickly and effectively.

Another limitation is access to the compute infrastructure required by AI research. Underserved communities may lack the financial resources to procure and manage appropriate infrastructure, thereby inequitably limiting access to developmental opportunities for budding researchers. NAIRR can partially overcome this limitation by providing and managing the infrastructure, and by hosting the AI research platform in a web-accessible location so that a more diverse range of researchers can access AI methods without needing their own infrastructure.

In summary, NAIRR can help democratize access to AI by developing a technology solution which:

- Prioritizes multi-model machine learning of various programming languages as well as non-coding approaches to accommodate the various skill sets and preferences of the research community.
- Takes advantage of various high performance computing techniques so that models can be processed quickly and that researchers have the ability to validate innovative ideas on sufficiently large datasets.
- Automates performance tracking and re-training of models to ensure optimal results as additional data are accumulated and keeps the focus on innovation.
- Includes or integrates with a model inventory to maintain documentation, versioning and model lineage within a governed platform.
- Provides built-in model interpretability capabilities to help researchers explain results and uncover bias.
- Incorporates centralized governance to control access to and usage of data and models across various stakeholders

Federal Register Notice 86 FR 46278, <https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Steve Xiao

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



RFI Response: National AI Research Resource

Frontier AI Development Based on Bionics or Biologically Inspired Engineering

Dr. X. Steve Xiao, Savannah River National Laboratory

The current National Artificial Intelligence Initiative Act of 2020 has general good coverage for National Artificial Intelligence Research Resource. To speed up the AI development in the US and enhance our international competitiveness, new avenue needs to be explored as frontier research. Suggest adding resources on the development of AI driven by bionics or biologically inspired engineering. Here we redefine AI as implementation of algorithm and logics Nature already engineered, into electronics. The intelligence developed from the nature includes decision making, sensing, control and performance like dogs, cats, insects, fishes and even plants, as well as human beings. We can learn not only from human's decision making, but also animals and insects. For example, a toddler learns to tell dogs from cats with only a few examples, while AI would need thousands of pictures to train. Dogs and cats walking elegantly are probably driven at top level by conservation of energy instead of coordination of joints and limbs. A tiny mosquito can sense CO₂ from long range, navigate towards potential pray, find the warm blood skin temperature, verify the host with skin chemical (octenol), suck blood, ride with wind back to the pond, and lay eggs to complete a life cycle. The Venus fly trap plant has simple nerves similar to animals. The neural response is verifiable with electrode. A spike is observed if trigger hair is touched, but the trap leaf is not ready to close. It turns out that the Venus fly trap can count one and two. The trap leaf closes if the trigger hairs are stimulated twice within 10 seconds, a simple logic to avoid false alarms from rain drops and dust fall. After about 30 second, the 1st stimulus is reset, and the plant starts over to count one and two. The simplicity and fool-proof decision making are what we could implement in the nowadays AI, and there are vast examples in the nature to copy with. This is a holy grail to be discovered. Nature neural networks have the computational power not in GHz, nor MHz, and likely in the kHz range or slower; however, nature creatures have been able to make correct decisions better than nowadays computers and AI programs in many areas, plus additional tasks on food hunting, reproduction, life cycle, etc. If the decision algorithm is understood and implemented into computers and electronics, it can be augmented by million times faster with GHz speed computer, or by million copies of the electronic devices. The biological neural networks are different from silicon-based gates and electronics. The frontier AI development based on bionics can be derived from the following directions:

- Understand the algorithm and logics nature engineered, based on bionics or biologically inspired engineering. Pay attention on their simplicity and fool-proof decision-making methodology as well as physiology differences from silicon-based electronic chips.
- Development of hybrid analog and digital computers. The biological nerves and signal processing are not entirely digital. They combine digital, analog, and chemical signals as harmonically orchestrate. Electronic analog signal processing is approximate but fast. Analog computer elements can be build based on

operational amplifiers. They perform addition/subtraction, differential, and integration easily. Arrays of analog processing elements can be arranged with digital gates to combine into complex and versatile tasks. The combination of analog signal processing with modern digital computation capability may be superior than digital alone.

- Holographic Modality Perception for Reinforcement Learning. Animals learn from nature's feedbacks with full spectrum of sensors. Example of humans' sensors includes vision, hearing, touching, temperature, moisture, taste, smell, acceleration, vibration, etc. This may be necessary to learn nature's intelligence as in the real world all sensors always agree each other. The modern AI ML (Machine Learning) commonly relies on a single data source, and occasionally combines two (e.g., audio, visual) as cross modality perception. A holographic modality perception may be necessary for advanced learning mimic nature.
- Identify correct feedbacks for the AI ML process. As the nowadays AI ML starts to have breakthroughs on successful tasks such as voice recognition and facial recognition, there are still immense areas to be explored. The job function of a future programmer could be similar to nanny and school teachers, who guide computers with feedbacks of awards and punishments while the computer AI system complete the program language behind the scenes.

In summary, a successful AI will learn and implement into electronics systems the algorithm and logics Nature already engineered. We may be able to avoid big data and HPC and work smart to achieve more advanced intelligence and decision-making by computer and electronics.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

**Abas Abdoli, Ryan N Coffee, Auralee
Edelen, Michael Kagan, Daniel Ratner,
Sohail Reddy, and Kazuhiro Tera**

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Request for Information Response:
National Artificial Intelligence Research Resources Task Force

Abas Abdoli, Ryan N Coffee, Auralee Edelen, Michael Kagan, Daniel Ratner, Sohail Reddy, and Kazuhiro Terao
SLAC National Accelerator Laboratory

Applications of artificial intelligence and machine learning (AI/ML) to the sciences are already ubiquitous but have not reached their full potential, with scalability and access holding back researchers across the spectrum. This response focuses on **Question 2, Topic D**, concerning opportunities and specific steps that will enable AI/ML methods to increase their impact on the sciences and support researchers throughout the US.

A common thread among the most well-known applications of AI/ML in industry is the enormous scale of resources required. As an example, the Generative Pre-trained Transformer-3 (GPT-3) contains many billions of parameters, processed nearly a trillion words, cost more than \$10M in compute, and required a dedicated team of ML and software engineers. Likewise in the sciences, one of the heralded AI/ML advances of the last year was the release of AlphaFold, again the product of a large, well-funded team from private industry with access to heavy resources in compute and engineering expertise and working with open scientific data. The scale of the AlphaFold model in combination with science domain knowledge proved a highly successful combination. While far smaller than GPT-3, AlphaFold is still larger than models commonly used in science today. While similar opportunities for AlphaFold-like models exist across the sciences, large-scale industry-built models are rarely directly applicable to scientific goals without significant modification and from-scratch retraining. It is vital to create ML/AI tools that are adapted to scientific data and tasks, and crucially are capable of scaling to the size where industrial AI/ML has seen break-throughs. Moreover, the real impact of AI/ML methods will appear when tools are available to all researchers, including students, individual contributors, and small academic groups across the breadth of research institutions in the US. Crucially, the tools must exist across the AI/ML lifecycle, including data collection, training, optimization, and deployment for automation and autonomous experiments. Widespread access to large AI/ML tools will both impact equity and innovation -- AI/ML tools should not be restricted to those lucky enough to have large resources -- and also will spread impact to scientists working on the most pressing scientific challenges of the day.

As scientists in the Department of Energy's national lab system, we see an opportunity to leverage existing investments in large-scale scientific experiments that can drive

AI/ML R&D broadly across the sciences. In particular, the DOE has invested in both computing resources and scientific facilities that generate the vast amounts of compute and data needed to train impactful AI/ML models. On the computing side, the current generation of leadership computing facilities at Argonne and Oak Ridge will push into the exascale while using AI/ML-friendly architectures. On the scientific-facility side, experiments are generating ever larger datasets. For example, the Vera Rubin Observatory will generate terabytes of data per day with its gigapixel camera, and SLAC's new x-ray laser facility will generate terabytes of data every second while operating around the clock. Even smaller scale instruments such as electron microscopes can produce 10 terabytes of images per day. The DOE's data and computing resources can be combined through the development of well-structured, benchmark datasets tied to specific science challenges, and paired with access to AI accelerating computational infrastructure like CPU clusters hosting farms of GPUs as well as emerging AI-tailored architectures such as developed by the likes of Google, Cerebras, Samba Nova, and GraphCore. Critically, access to datasets and such computing resources should be available to researchers across the spectrum of sciences and institutions.

Software Engineering: While computing and datasets are central to building AI/ML models, large-scale training also requires significant effort and expertise in software engineering. The requisite level of labor may not be available to individuals or small groups. We see several paths to solving this problem: creation of a software engineering workforce for AI/ML and science, investment in education, and development of open-access tools. At the national lab level, ML/software engineer teams could be leveraged across research groups to scale AI/ML projects. Democratizing engineering resources beyond national labs will be more challenging, but for example national engineering support for academic research could have a multiplicative effect on scientific progress. Any solution must be sufficiently flexible to support projects that vary from a small group of students to national collaborations with thousands of researchers. Investment in education is required both to create a dedicated workforce of software engineers as well as to educate scientists themselves. Researchers require training in both standard data science as well as in topics tailored to scientific AI/ML, such as uncertainty quantification, probabilistic models, and handling multi-modal data. Large-scale training -- requiring clusters and sophisticated data structures -- will be increasingly critical. Finally, open-access tools reduce the barrier to entry for novice AI/ML practitioners. Existing examples from industry such as PyTorch (from Facebook) and TensorFlow (from Google) have played an important role in the current democratization of AI/ML. Similarly, tools should be developed specifically for scientific challenges, for example handling new types of data, providing robust statistical analysis, or enabling large-scale training on clusters. Pre-trained models could also be

considered tools. For example, in the same way that the GPT-n family of pre-trained transformers was conceived as a generic tool for natural language processing, large-scale models pre-trained on domain-wide scientific datasets could be adapted to individual research tasks in new settings and with less labeled data in a sub-field with vastly reduced effort and investment.

Benchmark Datasets: The availability of benchmark datasets and associated models are key to democratizing AI/ML. However, shared infrastructure raises a host of issues to be addressed. In industrial applications, large-scale, public data (e.g. ImageNet, COCO, ShapeNet for Computer Vision) played a critical role in the breakthrough of modern AI/ML: benchmark datasets fuel the training of AI/ML models, enable transparent and fair comparison between solutions, and support interdisciplinary collaborations among domain experts. Access to benchmark datasets however remains a non-trivial challenge for most scientific datasets. Even within a single experimental collaboration, it may take years of training for new members to interpret data structures and develop tools necessary for a scalable AI/ML solution. Independent groups with no or limited access to the data and lacking knowledge of the data provenance are simply unable to participate. Furthermore, most tool development happens without a public or an inter-experimental collaboration in mind, thus requiring duplication of effort. An organized effort that encourages scientists to design databases and interfaces that enable efficient AI/ML development and sharing is urgently needed. Such a database should provide easy access (e.g. widely known data format, with community accepted visualization tools), detailed description of the contents (e.g. attribute descriptions, a complete metadata), definitions of scientific benchmark metric and test datasets, and direct coupling with data storage, distribution, and computing infrastructures behind the scene to allow efficient AI/ML model development. All of these components require support from software and data engineers as well as training resources for scientists to design effective solutions.

The goals enumerated here are challenging and the proposed solutions will require significant investment, but the potential for return on investment is enormous: new types of data analysis, improved and even autonomous operations and performance of instruments, and more. For the rest of this document, we give additional specific, actionable steps that relate to the topics in the RFI.

Question 1, Topic B, The appropriate agency responsible for the research resource:

It is recommended that the US Department of Energy (DOE) play a role in generating and hosting large-scale benchmark scientific datasets and associated AI models. Given

the number of scientific facilities operated by national laboratories paired with exascale computing resources, the DOE is in a prime position to contribute to the NAIRR effort. DOE facilities produce an overwhelming abundance of high-quality, multi-disciplinary datasets that span a wide range of scientific fields and span the spectrum from openly public to nationally secure datasets. These datasets include some of the largest scientific experiments operated by the federal government, as well as numerous smaller facilities in physics, chemistry, biology, materials, geology, environmental science, and more. On the computing side, the DOE is currently building the first two public exascale machines in the US -- Frontier at Oak Ridge National Laboratory and Aurora at Argonne National Laboratory -- both set to come online in 2022. In addition, the National Energy Research Scientific Computing Center (NERSC) is currently installing its own upgrade in Perlmutter. Predecessors (Summit and MIRA) may also be repurposed for NAIRR efforts, and additional computing systems, in particular IoT and Edge computing systems, exist throughout the DOE laboratory system. Due to the increasingly large size of scientific datasets -- commonly in the terabytes and reaching into the petabyte scale -- being generated in the DOE system, it is recommended that the computing resources be connected directly to the benchmark datasets and over network to the scientific instruments with the highest data production rates that are soon to exceed TB/second continuous operation.

Question 1, Topic C, Model for allocation of resources:

It is recommended that NAIRR initiative follow the tier-like allocation model currently used to allocate resources e.g. at ORNL on Summit. This model requires several computing architectures with variable levels of scalability and requires users to demonstrate the scalability/performance of their method/model on less intensive architecture before allocating high-tier resources. Such resource allocation may also be extended to ML engineering / software engineering support. Through an engineering workforce created through the NAIRR initiative and potentially in partnership with the private sector, users could request ML and software engineering support from a broad community. If a request meets the criteria set by the engineering review, it would be allocated engineering support. This engineering time could be requested on user grant applications. In addition to this model, it is recommended that the allocation of resources consider the following criteria:

- 1) Impact/significance of the research
- 2) Feasibility
- 3) Novelty
- 4) Benefits of research (in terms of public availability or inter-agency benefit of the resulting data/results, including equity)

- 5) Degree of collaboration across multiple institutions and/or federal agencies

Question 1, Topic D, Capabilities needed to create shared infrastructure

We have seen that both models and datasets can be cross-disciplinary, in the sense that a model or dataset from one domain can be adapted or applied in other domains. It is recommended that a set of procedures/guidelines be developed that dictate the structure and interfaces for both models and datasets. For example, foundational architectures (e.g. GPT for text generation) can be tuned to perform a wide range of specific niche tasks within text generation, but more importantly for science these architectures can be re-trained from scratch and applied to completely new domains (e.g. music). To increase reusability and portability of previously developed models, a set of standards should be developed for model storage, training, deployment, etc. Similarly, a centralized set of standards for benchmark datasets in scientific domains is needed to govern data storage formats, access, and metadata to reduce engineering overhead and lower the barrier to training and comparing model performance. The chosen standards should follow the Findability, Accessibility, Interoperability, and Reusability (FAIR) data principles. As datasets are domain-specific, it is recommended that datasets are created and provided by domain experts and guided by standards and engineering tools to encourage FAIR principles. Such datasets can then be made accessible through NAIRR with appropriate access control. We expect that few datasets would need to be stored at the NAIRR facilities, but rather would be provided for remote access under a NAIRR licensing agreement.

Question 1, Topics E,F,G, Question 5: Limitation of NAIRR ability and Federated Machine Learning

A limitation of the NAIRR is the ability to generate and distribute data. Although data in science is abundant, at present only the private sector has the required data to develop AI for applied everyday use. The data gathering practices at Facebook, Google, Amazon, etc. allow the private sector to both develop and deploy their AI for mainstream uses. The involvement of NAIRR in gathering certain types of scientific data under a federal agency poses both privacy concerns and potentially national security concerns. The DoE is well positioned to apply security and encryption tools to data in a federal-scale AI initiative, and may actually be an opportunity to develop AI solutions that are more broadly applicable to e.g. financial services industry and healthcare.

Concerns regarding data and model dissemination (topic E) are fundamentally linked to both security (topic F) and privacy (topic G), and are as yet unaddressed by the scientific ML community. A Federated Machine Learning (FedML) ecosystem supports training across distributed datasets without need for direct sharing of data and would enable collaboration even among competing institutions. FedML leverages diverse datasets to ensure large models generalize rather than shoehorn into narrow niche applications. Large pre-trained models such as Bidirectional Encoder Representations from Transformers (BERT) and GPT-style transformers must be pre-trained on a broad distribution of training examples. For these language models, a tremendously diverse set of examples are required such as all of Wikipedia text. The resulting models can then be applied to niche problems with scarce labels. However, in domains where data is compartmentalized, for example due to privacy or national security concerns, it is not possible to assemble a single dataset of sufficient size to train a foundational architecture.

A well constructed FedML ecosystem could alleviate this issue with e.g. homomorphic encryption based ML training whereby loss computation neither exposes the model weights nor data samples of the host repository. This could allow for even niche science cases to contribute to and benefit from a generalizable base model. Already, the financial services industry has begun to derive market value from FedML architectures in which competing institutions leverage shared knowledge while still preserving their individual competitive advantage. We feel that such a model for science would allow for inter-domain and even inter-lab competition while nevertheless encouraging an integrated web of robust and evolutionarily improving scientific base models without ever exposing sensitive or private data.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Stanford Libraries

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



Stanford | LIBRARIES

31 August 2021

RFI Response: National AI Research Resource, an RFI from the National Science Foundation and The Science and Technology Policy Office of the President of the United States of America

We believe an ecosystem for computing and data infrastructure for AI researchers and students should include librarians and library practices to do the critical work of acquisition (licensing, purchasing, accepting locally built or mashed-up data sets), curating, managing, preserving, and providing means of discovery for data, models, and any other residual outcomes of AI research.

This response addresses Component D of the roadmap: the capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.

1. What options should the Task Force consider and why?

In response to Component D, the Task Force should recognize that research libraries and archives are well-positioned to address the growing need for solutions for management and discovery for the large-scale datasets required for AI research. Libraries have for centuries been supporting the needs of research and have evolved in response to the changing needs of researchers. The AI community has struggled, and too often failed, to address concerns of privacy, protection of intellectual property, transparency, and democratization, all of which are core values that libraries and library practices have been developed to address. Indeed, many research libraries are already providing these services and support for data users at universities.

2. Which capabilities and services provided through the NAIRR should be prioritized?

A shared computing and data infrastructure should be built on the foundational principles of libraries, archives, and museums: that information must be not just preserved, but discoverable; not just discoverable, but deliverable; not just deliverable as bits, but readable; not just readable, but understandable; and not just understandable, but usable (OAI Reference Model <https://public.ccsds.org/Pubs/650x0m2.pdf>).

The specific capabilities to be prioritized include:

- Application Program Interfaces (APIs). This is an initial step in lowering the barrier to access; necessary but not sufficient.
- Data sub setter for search and selection. Provide interactive filtering and selection based on metadata and statistical measures of a dataset.
- Delivery of data at scale. A pipeline from query and selection (as in 1.2) to delivery.
- Roundtrip derivatives from Machine Learning work to digital preservation and make them available, in context, via the delivery platform.
- Text extraction. Expand existing processes like OCR to include Handwritten Text Recognition and speech to text. Just as machine learning based OCR has transformed discovery, similar techniques can make a range of other materials, including audio and video, available to content navigation.
- Vectorization/Feature extraction pipelines. Representing text and images numerically makes it possible to perform meaningful analytics on this derivative form of the work and is a necessary first step for applying machine learning algorithms.
- Use library experience with knowledge bases and authorities to extend the value of named entity extraction (NER). Connect entities stored in records across databases and archives, identifying and providing means for verification and disambiguation.
- Build on the library's investment in RDF data structures to manage objects and their relationships. Once a content object's entities or metadata descriptors have been linked to authorities and knowledge bases, this becomes a basis for linking content across collections.
- Make the machine-actionable data discoverable nationally based on adoption of the DDI model.
- Apply computational analysis to make the machine-actionable data discoverable. The RDF graph is easily integrated into a machine learning workflow to produce analytics and visualizations to support curation as well as discovery.
- Bring the library model of subject specialists to contextual data analysis. Subject specialists as curators can make use of new quantitative views on collections to better understand their characteristics, gaps, and distributions to serve researchers.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Business and research practices that prioritize optimization and efficiency over equity and a contextual, historical understanding of the data feeding the models are inherently discriminatory. The fact that AI, emerging as it has from computer science, lacks a grounding in ethical principles of data collection, management, and use should not be a surprise. Researchers across the natural sciences, social sciences and humanities rely on libraries to record provenance, provide context, make data sets discoverable, preserve, and perform maintenance, and manage data. While those actions are not neutral, they take place within

social and technical systems that take into account the changing and contested nature of information. They give researchers the tools to explore and interrogate questions of fairness, bias, and accountability.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

In the increasingly ethically challenged world of machine learning, libraries are not only ideally positioned but have an obligation to bring their considerable experience and expertise to bear as stewards of training data -- its creation, documentation, preservation, and reuse. Libraries have long been the centers for information and education in their communities, and those shared resources and the commitment to their care and preservation is a core value of what libraries do.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

The non-commercial service model of libraries encourages resource sharing and consortial agreements, making it possible for libraries to provide access to collections beyond what they already offer. Those offerings are wholly democratic to begin with, as they are a shared resource for all members of its community. Libraries do not seek profit as a core mission, and do not face the conflicts of interest that private businesses may have. Any member of the community of a library has access to it, and that access is far cheaper and wide-ranging than what most individuals could hope to attain for themselves. We aim to serve everyone for free and leverage collective resources to do so. Libraries also observe restrictions and limitations to use based on negotiated terms and government regulations, recording data use agreements to specific data sets when necessary.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Though libraries increasingly provide access to digitized content, systems are still, by and large, oriented to access to one object at a time or on a collection by collection basis. Platforms that can manage these data collections and their derivatives, link features across them, and provide tools for analysis, discovery, selection, and delivery are necessary. Shared collections repositories such as HathiTrust, which have provided access to collections through programs such the Emergency Temporary Access Service to enable print collections to be unlocked during COVID-19 lockdown restrictions, or who have provided hosting for groups like the Technical Report Archive and Image Library to digitize and make accessible federal government technical reports which might otherwise be lost or inaccessible, are key partners to libraries and can provide models for cross-object/cross-collection access.

Thank you for your attention to this response. We stand ready to discuss this response, which is submitted by the Stanford Libraries and is likely not the only response that will be submitted by agencies and individuals at Stanford.

Respectfully submitted,

Michael A. Keller
Vice Provost and University Librarian
Director of Academic Information Resources
Stanford University

Stanford Libraries:
101 Green Library
Stanford, CA 94305-6004
U.S.A.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Stanford University Institute for Human-Centered Artificial Intelligence (HAI)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



Stanford University
Human-Centered
Artificial Intelligence

October 1, 2021

Lynne Parker, Ph.D.
Director, National AI Initiative
White House Office of Science and Technology Policy

Erwin Gianchandani, Ph.D.
Senior Advisor, Translation, Innovation, and Partnerships
National Science Foundation

Dear Dr. Parker and Dr. Gianchandani,

We are pleased to submit the attached in response to the Request for Information by the White House Office of Science and Technology and National Science Foundation for creating and implementing a National Artificial Intelligence Research Resource (NAIRR).

Stanford's Institute for Human-Centered Artificial Intelligence (HAI) has long championed the creation of the NAIRR. In 2020, Stanford HAI led efforts with 22 top universities and a bipartisan, bicameral group of lawmakers to pass legislation establishing the NAIRR Task Force. In 2021, we convened a policy practicum course at Stanford on *Creating a National Research Cloud*, which brought together law, business, and engineering students, researchers, and faculty to address some of the core questions for building the NAIRR. To prepare the Report, we interviewed a wide range of stakeholders, engaged in extensive research on existing models, and developed a set of recommendations on both the data and compute dimensions.

To adhere to the 10-page limitation for submissions, please find enclosed the Executive Summary from our Report. Because our full Report, "Building a National AI Research Resource: A Blueprint for the National Research Cloud," spans 100+ pages, we have enclosed a draft as supplemental material. We anticipate that the final Report will be made publicly available the week of October 4.

We hope our extensive work on the design of the NAIRR proves useful to the Task Force and we are happy to engage the Task Force in any way we might be helpful.

Sincerely,

Daniel E. Ho, Ph.D.
William Benjamin Scott and Luna M. Scott
Professor of Law
Associate Director, Stanford HAI

Jennifer King, Ph.D.
Privacy and Data Policy Fellow, Stanford HAI

Russell C. Wald
Director of Policy, Stanford HAI

Christopher Wan
JD/MBA Candidate, Stanford University

Executive Summary: Creating a National Research Cloud¹

Artificial intelligence (AI) appears poised to transform the economy across sectors ranging from healthcare and finance, to retail and education. What some have coined the “Fourth Industrial Revolution”² is driven by three key trends: greater availability of data, increases in computing power, and improvements to algorithm design. First, increasingly large amounts of data have fueled the ability for computers to learn, such as by training an algorithmic language model on all of Wikipedia.³ Second, better computational capacity (often termed “compute”) and compute capability have enabled researchers to build models that were unimaginable merely 10 years ago, sometimes spanning billions of parameters (an exponential increase in scope from previous models).⁴ Third, basic innovations in algorithms are helping scientists to drive forward AI, such as the reinforcement learning techniques that enabled a computer to defeat the world champion in the board game Go.⁵

Despite these trends, the AI innovation landscape faces serious potential challenges. Historically, partnerships between government(s), universities, and industries have anchored the U.S. innovation ecosystem. The federal government played a critical role in subsidizing basic research, enabling universities to undertake high-risk research that can take decades to commercialize. This approach catalyzed radar technology, the internet, and GPS devices. As the economists Ben Jones and Larry Summers put it, “[e]ven under very conservative assumptions, it is difficult to find an average return below \$4 per \$1 spent” on innovation, and the social returns might be closer to \$20 for every dollar spent.⁶ Industry in turn, scales and commercializes applications.

Core challenges to this ecosystem and the future of AI exist. Computing power has become critical for the advancement of AI, but the high cost of compute has placed cutting edge AI research in a position accessible only to key industry players and a handful of elite universities.⁷ Access to data—the raw ingredients used to train most AI models—is increasingly

¹ This is the Executive Summary to a broader report on the National Research Cloud: DANIEL. E. HO, JENNIFER KING, RUSSELL C. WALD, CHRISTOPHER WAN, BUILDING A NATIONAL AI RESEARCH RESOURCE: A BLUEPRINT FOR THE NATIONAL RESEARCH CLOUD (2021).

² KLAUS SCHWAB, THE FOURTH INDUSTRIAL REVOLUTION (2016).

³ Tae Yano & Moonyoung Kang, *Taking Advantage of Wikipedia in Natural Language Processing*, CARNEGIE MELLON U. (2008), <https://www.cs.cmu.edu/~taey/pub/wiki.pdf>.

⁴ See, e.g., Anthony Alford, *Google Trains Two Billion Parameter AI Vision Model*, INFOQ (June 22, 2021), <https://www.infoq.com/news/2021/06/google-vision-transformer/>; Anthony Alford, *OpenAI Announces GPT-3 AI Language Model with 175 Billion Parameters*, INFOQ (June 2, 2020), <https://www.infoq.com/news/2020/06/openai-gpt3-language-model/>.

⁵ *AlphaGo*, DEEPMIND (2021), <https://deepmind.com/research/case-studies/alphago-the-story-so-far/>.

⁶ Benjamin F. Jones & Lawrence H. Summers, *A Calculation of the Social Returns to Innovation* (Nat’l Bureau of Econ. Research, Working Paper No. 27863, 2020); J.G. Tewksbury, M.S. Crandall & W.E. Crane, *Measuring the Societal Benefits of Innovation*, 209 SCI. MAG. 658-62 (1980); see also NATIONAL ACADEMIES OF SCIENCES, ENGINEERING, AND MEDICINE, RETURNS TO FEDERAL INVESTMENTS IN THE INNOVATION SYSTEM (2017)

⁷ STUART ZWEBEN & BETSY BIZOT, 2019 TAULBEE SURVEY: TOTAL UNDERGRAD CS ENROLLMENT RISES AGAIN, BUT WITH FEWER NEW MAJORS; DOCTORAL DEGREE PRODUCTION RECOVERS FROM LAST YEAR’S DIP (2019).

limited to the private sector and large platforms⁸, since government data sources remain largely inaccessible to the AI research community.⁹ As the National Security Commission on AI (NSCAI) has determined, “[t]he consolidation of the AI industry threatens U.S. technological competitiveness.”¹⁰ Four interrelated challenges illustrate this finding: First, we are seeing a significant brain drain of researchers departing universities.¹¹ In 2011, AI Ph.D.s were roughly equally likely to go into industry vs. academia.¹² Ten years later, two-thirds of AI Ph.D.s go to industry, and less than one quarter go into academia.¹³ Second, these trends indicate that many university researchers struggle to engage in cutting-edge science, draining the field of the diverse set of research voices that it needs. Third, the fundamental research that would guarantee the United States stays at the helm of AI innovation is being crowded out. By one estimate, 82 percent of algorithms used today originated from federally funded nonprofits and universities, but “U.S. leadership has faded in recent decades.”¹⁴ Fourth, government agencies have faced challenges in building compute infrastructure,¹⁵ and there are societal benefits to reducing the cost of core governance functions and improving government’s internal capacity to develop, test, and hold AI systems accountable.¹⁶ In short, a growing imbalance in AI innovation tilts towards industry, leaving academic and non-commercial research behind. Given the longstanding role of academic and non-commercial research in innovation, this shift has substantial negative consequences for the American research ecosystem.

Responding to these challenges, Congress enacted the National AI Research Resource Task Force Act as part of the National Defense Authorization Act (NDAA) in January 2021.¹⁷ The Act forms part of the National Artificial Intelligence Initiative, which identifies further steps to increase research investments, set technical standards, and build a stronger AI workforce. The Act created a Task Force—the composition of which was announced on June 10, 2021¹⁸—to study and plan for the implementation of a “National Artificial Intelligence Research Resource,” namely “a system that provides researchers and students across scientific fields and disciplines with access to compute resources, co-located with publicly available, artificial intelligence-ready

⁸ Jathan Sadowski, *When Data is Capital: Datafication, Accumulation, and Extraction*, 2019 BIG DATA & SOC’Y 1 (2019).

⁹ Amy O’Hara & Carla Medalia, *Data Sharing in the Federal Statistical System: Impediments and Possibilities*, 675 ANNALS AM. ACAD. POL. & SOC. SCI. 138, 140-41 (2018).

¹⁰ NAT’L SECURITY COMM’N ON ARTIFICIAL INTELLIGENCE, FINAL REPORT 186 (2021).

¹¹ STAN. U. INST. FOR HUMAN-CENTERED ARTIFICIAL INTELLIGENCE, 2021 ARTIFICIAL INTELLIGENCE INDEX REPORT 118 (2021).

¹² *Id.*

¹³ *Id.*

¹⁴ Neil C. Thompson, Shuning Ge & Yash M. Sherry, *Building the Algorithm Commons: Who Discovered the Algorithms that Underpin Computing in the Modern Enterprise?*, 11 GLOBAL STRATEGY J. 17-33 (2020).

¹⁵ See, e.g., U.S. GOV’T ACCOUNTABILITY OFFICE, FEDERAL AGENCIES NEED TO ADDRESS AGING LEGACY SYSTEMS (2016); U.S. GOV’T ACCOUNTABILITY OFFICE, CLOUD COMPUTING: AGENCIES HAVE INCREASED USAGE AND REALIZED BENEFITS, BUT COST AND SAVINGS DATA NEED TO BE BETTER TRACKED (2019).

¹⁶ DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES 6, 71-72 (2020).

¹⁷ William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, § 5106.

¹⁸ *The Biden Administration Launches the National Artificial Intelligence Research Resource Task Force*, THE WHITE HOUSE (June 10, 2021), <https://www.whitehouse.gov/ostp/news-updates/2021/06/10/the-biden-administration-launches-the-national-artificial-intelligence-research-resource-task-force/>.

government and non-government data sets.”¹⁹ This research resource has also been referred to as the National Research Cloud (NRC) and was strongly endorsed by the NSCAI, which wrote that the NRC “will strengthen the foundation of American AI innovation by supporting more equitable growth of the field, expanding AI expertise across the country, and applying AI to a broader range of fields.”²⁰

While other initiatives have sought to improve access to compute or data in isolation,²¹ the NRC will generate distinct positive externalities by integrating compute and data, the two bottlenecks for high-quality AI research. Specifically, the NRC will provide affordable access to high-end computational resources, large-scale government datasets in a secure cloud environment, and the necessary expertise to benefit from this resource through a close partnership between academia, government, and industry. By expanding access to these critical resources in AI research, the NRC will support basic scientific AI research, the democratization of AI innovation, and the promotion of U.S. leadership in AI.

Stanford Law School’s policy practicum program convened a multidisciplinary research team of graduate students, staff, and faculty drawn from Stanford’s business, law, and engineering schools to study the feasibility of, and considerations for designing the NRC. Over the past six months, this group studied existing models for compute resources and government data, interviewed a wide range of government, computer science, and policy experts, and examined the technical, business, legal, and policy requirements. This report was commissioned by Stanford’s Institute for Human-Centered Artificial Intelligence (HAI), which originated the proposal for the NRC in partnership with 21 other research universities.²²

Throughout our research, we observed three primary themes that cut across all areas of our investigation. We have integrated these themes into each section of our report and drawn on them to explain our findings.

- *Complementarity between compute and data.* As we evaluated the existing computing and data-sharing ecosystems, one of the systemic challenges we observed was a decoupling of compute resources from data infrastructures. High-performance computing can be useless without data; and a major impediment to data sharing, particularly for high-value government data, lies in requirements for a secure, privacy-protecting computing environment.
- *Rebalancing AI research toward long-term, academic, and non-commercial research.* Presently, AI innovation is disproportionately dependent on the private sector. Public

¹⁹ William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, § 5107 (g).

²⁰ NAT’L SECURITY COMM’N ON ARTIFICIAL INTELLIGENCE, *supra* note 10, at 191.

²¹ See, e.g., *Cloudbank*, <https://www.cloudbank.org>; *Fact Sheet: National Secure Data Service Act Advances Responsible Data Sharing in Government*, DATA COALITION (May 13, 2021), <https://www.datacoalition.org/fact-sheet-national-secure-data-service-act-advances-responsible-data-sharing-in-government/>.

²² Steve Lohr, *Universities and Tech Giants Back National Cloud Computing Project*, N.Y. TIMES (June 30, 2020), <https://www.nytimes.com/2020/06/30/technology/national-cloud-computing-project.html>; John Etchemendy & Fei-Fei Li, *National Research Cloud: Ensuring the Continuation of American Innovation*, STAN. U. INST. FOR HUMAN-CENTERED ARTIFICIAL INTELLIGENCE, (Mar. 28, 2020), <https://hai.stanford.edu/news/national-research-cloud-ensuring-continuation-american-innovation>.

investment in basic AI infrastructure can both support innovation in the public interest and complement private innovation efforts. The NRC directs more resources toward AI development in the public interest and helps ensure long-term leadership by the United States in the field by supporting the kind of pure, basic research that the private sector cannot undertake alone.

- *Coordinating short-term and long-term approaches to creating the NRC.* Our research considers many near-term pathways for standing up a working version of the NRC by spelling out how to work within existing constraints. We also identify the structural, legal, and policy challenges to be addressed in the long term for executing the full vision of the NRC.

We summarize our main recommendations here.

Compute Model

- The “Make or Buy” Decision. The main policy choice will be whether to build public computing infrastructure or purchase services from existing commercial cloud providers.
 - It is well established that, based solely on hardware costs, it is more cost-effective to own infrastructure when computing demand is close to continuous.²³ The government also has experience building high-performance computing clusters, typically built by contractors and operated by national laboratories.²⁴ The National Science Foundation (NSF) has also supported many supercomputing initiatives at academic institutions.²⁵
 - The main countervailing concerns are that existing commercial cloud providers have software stacks and usability that AI researchers have widely adopted and may consider to be a more user-friendly platform. Commercial cloud providers offer a way to expand capacity expeditiously, although scale and availability will still be constrained by the availability of current graphics processing unit (GPU) computing resources.
 - We hence recommend a dual investment strategy:
 - First, the compute model of the NRC can be quickly launched by subsidizing and negotiating cloud computing for AI researchers with existing vendors, expanding on existing initiatives like the NSF’s CloudBank project.²⁶
 - Second, the NRC should invest in a pilot for public infrastructure to assess the ability to provide similar resources in the long run. Such publicly owned infrastructure would still be built under contract or grant, but could

²³ Jennifer Villa & Dave Troiano, *Choosing Your Deep Learning Infrastructure; The Cloud vs. On-Prem Debate*, DETERMINED AI (July 30, 2020), <https://determined.ai/blog/cloud-v-onprem/>; *Is HPC Going to Cost Me a Fortune?*, INSIDEHPC (last visited July 23, 2021), <https://insidehpc.com/hpc-basic-training/is-hpc-going-to-cost-me-a-fortune/>.

²⁴ See, e.g., *US Plans \$1.8 Billion Spend on DOE Exascale Supercomputing*, HPCWIRE (Apr. 11, 2018), <https://www.hpcwire.com/2018/04/11/us-plans-1-8-billion-spend-on-doe-exascale-supercomputing/>; *Federal Government, ADVANCED HPC* (last visited July 23, 2021), <https://www.advancedhpc.com/pages/federal-government>; *United States Continues to Lead World In Supercomputing*, U.S. DEP’T. ENERGY (Nov. 18, 2019), <https://www.energy.gov/articles/united-states-continues-lead-world-supercomputing>.

²⁵ See *NSF Funds Five New XSEDE-Allocated Systems*, NAT’L SCI. FOUND. (Aug. 10, 2020), <https://www.xsede.org/-/nsf-funds-five-new-xsede-allocated-systems>.

²⁶ *Cloudbank*, *supra* note 21.

be operated much like national laboratories (e.g., Sandia National Laboratories, Oak Ridge National Laboratory), that own sophisticated supercomputing facilities or academic supercomputing facilities.

- **Researcher Eligibility.** While some have argued the NRC should be open for commercial access, for the purposes of this report, we adhered to the spirit of the legislation forming the NAIRR Task Force and only reviewed the use of an NRC for academic and non-profit AI research. We recommend that the NRC eligibility start with academics who hold “Principal Investigator” (PI) status (i.e., most faculty) at U.S. colleges and universities, as well as to “Affiliated Government Agencies” willing to contribute previously unreleased, high-value datasets to the NRC in return for subsidized compute resources. PI status should be interpreted expansively to encompass all fields of AI application. Students working with PIs should presumptively gain access to the NRC. Scaling the NRC to meet the demand of all students in the United States may be challenging, but we also recommend the creation of educational programs as part of the new resource to help train the next generation of AI researchers.
- **Mechanism.** In order to keep the award processing costs down, we recommend a base-level of compute access to meet the majority of researcher computing needs. Base-level access avoids high overhead for grant administration and may meet the compute demands for the supermajority of researchers. For researchers with exceptional needs, we recommend a streamlined grant process for additional compute access.

Data Access Model

- **Focus on Government Data.** We focus our recommendations for data provision/access to government data because: (1) there are already a wide range of platforms for sharing private data,²⁷ and (2) distribution by the NRC of private datasets would raise a tangle of thorny IP issues. We recommend that researchers be allowed to compute on any datasets they themselves contribute, provided they certify they have the rights to that data, and the use of such data is for academic research purposes.
- **Tiered Access.** We recommend a tiered access model: by default, researchers will gain access to government data that is already public; researchers can then apply through a streamlined process to gain access at higher security levels on a project-specific basis. It will be critical for the NRC to ultimately displace the current fragmented, agency-by-agency relational approach. By providing secure virtual environments and harmonizing security standards (e.g., Federal Risk and Authorization Management Program (FedRAMP)²⁸), the NRC can collaborate with proposals for a National Secure Data Service²⁹ to provide a model for accelerating AI research, while protecting data privacy and prioritizing data security.
- **Agency Incentives.** To incentivize federal agencies to share data with the NRC and improve the state of public sector technology, we recommend the NRC permit federal

²⁷ See, e.g., *National Data Service*, <http://www.nationaldataservice.org>; *The Open Science Data Cloud*, <https://www.opensciencedatacloud.org>; *Harvard Dataverse*, <https://dataverse.harvard.edu>; *FigShare*, <https://figshare.com>.

²⁸ *FedRAMP*, <https://www.fedramp.gov>.

²⁹ See *Fact Sheet: National Secure Data Service Act Advances Responsible Data Sharing in Government*, DATA COALITION (May 13, 2021), <https://www.datacoalition.org/fact-sheet-national-secure-data-service-act-advances-responsible-data-sharing-in-government/>.

agency staff to use the NRC's compute resources. In keeping with the practices of existing data-sharing programs, such as the Coleridge Initiative,³⁰ we also recommend that the NRC provide training and support to work with agencies to modernize and harmonize their data standards.

- **Strategic Investment for Data Sources.** In the short term, we recommend that the NRC focus its efforts on making available non-sensitive, low- to moderate-risk government datasets, rather than sensitive government data (e.g., data about individuals) or data from the private sector, due to data privacy and intellectual property concerns. Researchers can still use NRC compute resources on private data, but should rely on existing mechanisms to acquire data for their own private buckets on the NRC. For example, images taken from Earth observation satellites, such as Landsat imagery, provide a promising low-risk, high-reward government dataset, as making such satellite imagery freely available to researchers has generated an estimated \$3-4 billion in annual economic benefits, particularly when combined with high-performance computing.³¹ Agencies such as the National Oceanic and Atmospheric Administration, the U.S. Geological Survey, the Census Bureau, the Administrative Office of the U.S. Courts, and the Bureau of Labor Statistics, for instance, also have rich datasets that can more readily be deployed. In the long run, access to high-risk datasets, such as those owned by the Internal Revenue Service (IRS) and the Veterans Administration (VA), will depend on the tiered access model.

Organizational Form

Where to institutionally locate the NRC poses a tradeoff between ease of coordination to obtain compute and ease of data access. For instance, locating the NRC within a single agency would make coordination with compute providers easier, but would make data access across agencies more difficult, absent further statutory authority. Many efforts to make data access to government data easier, most notably the Foundations for Evidence-Based Policymaking Act of 2018, have proven to be among the most daunting challenges of government modernization.³² Building on those insights, we ultimately recommend that the NRC be instituted as a Federally Funded Research and Development Center (FFRDC) in the short run, and a public-private partnership (PPP) in the long run.

- **FFRDC.** FFRDCs at Affiliated Government Agencies would reduce the significant costs of securing data from those host agencies. This approach will also cohere with the greater reliance on commercial cloud credits in the short run, making compute and data

³⁰ See *Administrative Data Research Facility*, COLERIDGE INITIATIVE, <https://coleridgeinitiative.org/adrf/> (last visited July 26, 2021).

³¹ See *Landsat Data Access*, U.S. GEOLOGICAL SURVEY, <https://www.usgs.gov/core-science-systems/nli/landsat/landsat-data-access> (last visited July 23, 2021); FED. GEOGRAPHIC DATA COMM., *THE VALUE PROPOSITION FOR LANDSAT APPLICATIONS* (2014); CRISTA L. STRAUB, STEPHEN R. KOONTZ & JOHN B. LOOMIS, *ECONOMIC VALUATION OF LANDSAT IMAGERY* (2019).

³² See BIPARTISAN POL'Y CTR., *BARRIERS TO USING GOVERNMENT DATA: EXTENDED ANALYSIS OF THE U.S. COMMISSION ON EVIDENCE-BASED POLICYMAKING'S SURVEY OF FEDERAL AGENCIES AND OFFICES 18-20* (2018); see also U.S. DEP'T OF HEALTH & HUMAN SERVICES, *THE STATE OF DATA SHARING AT THE U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES 4* (2018) (describing how data at the agency is "largely kept in silos with a lack of organizational awareness of what data are collected across the Department and how to request access.").

coordination less central. In the long run, however, streamlined coordination between data and compute may be more difficult with FFRDCs hosted at specific agencies when (1) the NRC moves away from commercial cloud credits and towards its own high-performance computing cluster, and (2) a greater number of inter-agency datasets become available.

- **PPP.** In the long run, we recommend the creation of a PPP model, governed by officers from Affiliated Government Agencies, academic researchers, and representatives from the technology sector, which can house both compute and data resources.

Additional Considerations

- **Data Privacy.** As an initial matter, an NRC where sensitive or individually identifiable administrative data from multiple agencies are used to build and train AI models will face challenges from the Privacy Act of 1974.³³ The Act is intended to put a check on interagency data sharing and disclosure of sensitive data without consent.
 - In order to avoid conflicts with non-consensual interagency data sharing, we recommend that the NRC should not be instituted as its own federal agency, nor should federal agency staff be allowed access to interagency data.
 - To avoid conflicts with the Act’s “no disclosure without consent” requirement, any data released to the NRC must not be individually identifiable. Despite these constraints, the majority of AI research will likely fall under the Act’s statistical research exception, contingent on proposals aligning with an agency’s core purpose.
 - Given concerns about the potential privacy risks, federal agencies may desire to share data, contingent on the use of technical privacy measures (e.g., differential privacy). While useful in many instances, technical approaches are no panacea and should not substitute for data access policies.
 - The NRC should explore the design of virtual “data safe rooms” that enable researchers to access data in a secure, monitored, and cloud-based environment.
 - Additional legislative interventions could also facilitate data sharing with the NRC (e.g., requiring IT modernization to include data sharing plans with the NRC).
- **Ethics.** Rapid innovation in AI research raises a host of potential ethical challenges. Given the scope of the NRC, it will be infeasible to review every single research proposal for potential ethical violations, particularly since ethical standards are still in flux. The NRC should adopt a twofold approach.
 - First, for default PI access to base-level data and compute, the NRC should establish an ex-post review process for allegations of ethical research violations. Access may be revoked when research is shown to manifestly and seriously violate ethical standards. We emphasize that the high standard for a violation should be informed by the academic speech implications and potential political consequences of government involvement in administering the NRC and determining academic research directions.
 - Second, for applications requesting access to restricted datasets or resources beyond default compute, which will necessarily undergo some review, researchers

³³ Privacy Act, 5 U.S.C. § 552a (1974).

should be required to provide an ethics impact statement. One of the advantages of beginning with PIs is that university faculty are accountable under existing IRBs for human subjects research, as well as to the tenets of peer review.

- We urge non-NRC parties (e.g., universities) to explore a range of measures to address ethical concerns in AI compute (e.g., an ethics review process³⁴ or embedding ethicists in projects³⁵).
- Security. We recommend that the NRC take the lead in setting security classifications and protocols, in part to counteract a balkanized security system across federal agencies that would stymie the ability to host datasets. The NRC should use dedicated security staff to work with Affiliated Government Agencies and university representatives to harmonize and modernize agency security standards.
- Intellectual Property (IP). While the evidence on optimal IP incentives for innovation is mixed, we recommend that the NRC adopt the same approach to allocating patent rights, copyrights, and data rights to NRC users that apply to federal funding agreements. The NRC should additionally consider conditions for requiring NRC researchers to disclose or share their research outputs under an open-access license.
- Human Resources. Given its ambition, significant human resources – from systems engineers to data officers, and from grants administrators to privacy, ethics, and cybersecurity staff – will be necessary to make the NRC a success.

³⁴ Michael S. Bernstein et al., *ESR: Ethics and Society Review of Artificial Intelligence Research*, CORNELL. U. (July 9, 2021), <https://arxiv.org/pdf/2106.11521.pdf>.

³⁵ Courtenay R. Bruce et al., *An Embedded Model for Ethics Consultation: Characteristics, Outcomes, and Challenges*, 5 *AJOB EMPIRICAL BIOETHICS* 8 (2014).

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

U.S. Chamber of Commerce Technology Engagement Center

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



MICHAEL RICHARDS

Director

1615 H STREET, NW
WASHINGTON, DC 20062-2000

October 1, 2021

The Office of Science and Technology Policy &
The National Science Foundation
Attn: Wendy Wigon, NCO
Alexandria, VA 22314

Re: Request for Information on an Implementation Plan for a National Artificial Intelligence Research Resource; 86 FR 46278

To Whom It May Concern:

The U.S. Chamber of Commerce's Technology Engagement Center ("C_TEC") appreciates the opportunity to submit comment to the Office of Science and Technology Policy (OSTP) and the National Science Foundation (NSF) Request for Information (RFI) on "an Implementation Plan for a National Artificial Intelligence Research Resource." C_TEC supports OSTP's and NSF's work to develop a National Artificial Intelligence Research Resource (NAIRR), which will produce a roadmap and implementation plan to create a shared computing resource for AI.

One of the guiding AI principles for C_TEC is to "promote Open and Accessible Government Data." C_TEC acknowledges that National Artificial Intelligence Research Resource will look to accomplish just that by leveraging large and robust government data sets, which will help spur further innovation and breakthroughs within the scientific community. We support NAIRR's commitment to making this data available in an accessible manner for the research community.

C_TEC wishes to provide the below feedback on OSTP's and NSF's request for information on the "Implementation Plan for a National Artificial Intelligence Research Resource." These comments include:

1. What options should the Task Force consider for roadmap elements A through I, and why?

a. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success;

C_TEC believes that setting goals and metrics are significant for tracking the success of the National Artificial Intelligence Research Resource. We agree that the following metrics and objectives should be considered: research usage, community uptake, and wider impacts.

Research Usage: One of the easiest ways to determine if research resources are being utilized is the overall usage. Therefore C_TEC believes that as NAIRR should use the metric of overall researcher usage.

Community Development: A critical metric to consider is the overall development of the U.S. AI community. Specifically, NAIRR should look at broad participation in workshops related to Artificial Intelligence, investment in R&D, the development of AI tools, the creation of new benchmarks or standards adopted by the community, and the number of experiments shared within the community.

Wider impacts: NAIRR will have the profound ability to assist in the development of resources that can drive innovation and lead the development of significant breakthroughs within scientific fields. This is why C_TEC believes that tracking breakthroughs within the field should be used as a metric, increases in productivity, automation, as well as the number of new products to enter into the market, and the number of startups created.

b. A plan for ownership and administration on the National Artificial Intelligence Research Resource, including:

i. An appropriate agency or organization responsible for the implementation, deployment, and administration of the Research Resource;

C_TEC can foresee multiple models for government ownership for the initial implementation, deployment, and administration phase for NAIRR. Whichever model is chosen, however, we encourage that it undergo review once the research resource has matured to ensure maximized utilization of the resource.

ii. A governance structure for the Research Resource, including oversight and decision making authorities;

C_TEC understands that NAIRR's effort to grow computing resources and high-quality data sets simultaneously is no easy task. For this reason, we encourage NAIRR to review the merits of housing the research resource within a Federally Funded Research and Development Center (FFRDC). FFRDC has a strong track record of spurring research and development and would allow multiple stakeholders to work together on the underlying goal of improving access to data to accelerate scientific discovery. Furthermore, C_TEC believes that multiple agency governance structures could be beneficial, in that it would allow for further buy-in from other agencies. We encourage the FFRDC to be connected with the Department of Energy (DOE), the National Institutes of Health (NIH), and other agencies, with input from OSTP.

c. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources;

C_TEC strongly supports allowing the industry to compete for the development and buildout of computing and data resources, AI tools, and platforms. We believe that the procurement of AI software and data management tools should be done in an open, transparent process. Furthermore, we would highlight the need for resources to be interoperable. However, we have concerns regarding total

homogenization of the resources as diversity ensures that researchers access the resources across multiple clouds and interfaces. Diversity is critical in allowing for quick utilization of the resources by researchers.

d. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provisions of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure;

Regarding data sets and secure access control, C_TEC supports the principles of findability, accessibility, interoperability, and reuse of digital assets or better known as the FAIR guiding principle for scientific data management and stewardship. Furthermore, we believe data sets should be easily located and legible for both human and machine consumption.

Regarding compute resources and scalability, C_TEC believes that NAIRR should take advantage of commercially available computing resources. The research community is already familiar with these resources and would allow for quick adoption and utilization. Furthermore, we believe NAIRR should adopt a hybrid, secure, multi-cloud approach to provide cost-effective computing at the necessary scale. NAIRR should also look for opportunities to leverage existing public clouds and security layers and integrate those in the multi-cloud.

Finally, regarding Educational Tools & Services/Resident Expertise, C_TEC believes that effective communication between researchers can spur collaboration and help further drive innovation and determine the reproducibility of results. We encourage collaboration between individual researchers, government entities, and industry and we support engagement necessary to exchange technical expertise and best practices in order to find synergies within their respective work.

e. An assessment of and recommended solutions to barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource;

C_TEC understands that many issues will need to be resolved before the dissemination of data sets by NAIRR. The following are areas we believe should be further addressed.

First, with the large amount of data coming from multiple agencies being compiled together, there are concerns that such data could have personally identifiable information (PII), Personal Health Information (PHI), or other sensitive data associated with it. For this reason, we encourage NAIRR to deidentify and secure all sensitive data. Also, NAIRR should look at how it can automate the ability to remove and protect sensitive data, as well as anonymize the data.

Second, the mass compiling of datasets could lead to issues with finding the correct data for specific research. C_TEC believes that it would be prudent to develop a hybrid data fabric to assist researchers in finding appropriate data more efficiently. Also, C_TEC would encourage the use of automated tagging and labeling, including looking into opportunities for supervised and unsupervised learning for tagging, to make data search and data understanding more accessible for researchers.

f. An assessment of security requirements associated with National Artificial Intelligence Research Resource and its management of access controls;

C_TEC believes that NAIRR will need to be secured to ensure that the data is not abused, misused, or otherwise compromised. We would encourage the use of security control such as access control, encryption, authentication, logging, and many others that may provide necessary security to the resources.

g. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence research resource and its research

Data is critical to the development of AI, and the repurposing of personal data may impact consumer privacy and trust in these research efforts. Therefore, clear and consistent privacy protections for individual privacy are necessary. However, should that not be plausible, we encourage NAIRR to implement robust but flexible data protection regimes that enable data collection, retention, and processing for AI development, deployment, and use while ensuring alignment with existing privacy laws.

2. What capabilities and services provided through NAIRR should be prioritized?

C_TEC believes multiple aspects of NAIRR should be prioritized simultaneously. These include the buildout of a hybrid cloud platform, the development of AI and data management software, open standards and open technologies, and consistent methods and tools.

First, NAIRR should prioritize developing a hybrid cloud platform that can provide a seamless user experience across multiple clouds.

Second, NAIRR should further prioritize the development of software that can assist the research resource's productivity, simplicity, portability, and reproducibility.

Third, NAIRR should prioritize the development of standards and open technologies which can assist with the adoption of cloud use and assist in reproducibility.

Finally, NAIRR should prioritize the development of consistent methods and tools to secure the data and resources.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as racial and gender equity, fairness, bias, civil rights?

C_TEC supports the focus on fairness and non-discrimination within NAIRR RFI. Fairness and non-discrimination principles are essential for establishing public trust in AI.

4. What building blocks are already for NAIRR regarding government, academic, or private sector activities, resources, and services?

C_TEC believes it is essential to review and incorporate other building blocks from existing collaborations to help NAIRR. Public initiatives such as the European Open Science Cloud, NIH STRIDES, and NSF Cloudbank have all connected researchers to datasets through the cloud. Other efforts that should be considered include the COVID-19 High-Performance Computing Consortium,

which was able to bring together academia, industry, and the government to harness high-performance computing to support COVID-19 research. Finally, we would highlight the Helix Nebula Science Cloud pilot program, creating a shared research space for millions of researchers.

5. What role should the public-private partnership play in the NAIRR? What examples could be used as a model?

C_TEC emphasizes that public trust-building efforts are conducted with government, industry, and other relevant stakeholders. A public-private partnership model will facilitate collaboration between all relevant stakeholders and allow for sharing of best practices.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And could these limitations be overcome?

NAIRR could experience many limitations to democratize the access to AI R & R&D, including a lack of the necessary workforce and skills needed to utilize the resources associated with research funding.

C_TEC believes in the use of commercially available resources. However, some researchers may face the challenges of learning how to use the resources and utilize them to accelerate their research fully. Therefore, an open line of communication between government, industry, and academia to learn best practices and necessary training to reduce the skills gap may be required.

NAIRR will provide further access to the necessary compute and data to allow for breakthroughs in scientific research. Yet, this does not solve the problem of having critical AI scientists and the research and development dollars to assist with their research. As the United States looks to develop this resource, we must also simultaneously look at ways to increase AI research transpiring in the field. This should include looking at federal research grants that may not have previously been used for AI research and developing the AI research workforce.

Conclusion

C_TEC appreciates NSF's and OSTP's ongoing work to develop NAIRR. We encourage further collaboration with stakeholders as we believe the partnership is vital for developing research resources that are fundamental for future scientific discovery. We thank you for your consideration of these comments and would be happy to further discuss any of these issues.

Sincerely,

Michael Richards
Director, Policy
Chamber Technology Engagement Center

Federal Register Notice 86 FR 46278, <https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence>, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

University of Florida

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

University of Florida comments

Version 3 9/16/2021 [Note: 10-page limit]

Contact: Erik Deumens

Given the partnership between UF and NVIDIA enabled by the donation of the fastest AI supercomputer in academia, called HiPerGator AI, given the vision of UF to create an AI University that will educate all students in all degree programs in the basics, the applications, the risks, the promises of AI, given that UF makes HiPerGator AI available to faculty and students in the State of Florida under the same policies as it is available to UF faculty and students, given that UF makes HiPerGator AI available to institutions in the SEC and select national universities for the purpose of teaching classes, UF was lead to consider ways to support a large and diverse user community with staff resources sized to support a single university.

These comments will address questions 1. Options, 2. Capability prioritization, and 5. Public-private partnerships.

Question 1. What options should the Task Force consider for any of roadmap elements A through I above, and why?

The comment addresses **B. Plan for ownership and administration i. Agency ii. Governance structure.**

Use a distributed model that is based on partnership with each institution willing to use the NAIRR to allow scalable support for all, not just the institutions and faculty who have sufficient expertise to use remote resources with minimal assistance.

Context: UF has 8 years' experience operating a large computing facility for its faculty and their students and national and international collaborators with a model that has proven to be sustainable. With the donation of HiPerGator AI, the Nvidia SuperPOD, to support UF's AI initiative, UF committed to provide access to a much larger user community, namely the faculty and their students and collaborators at all universities in the state. We set up a collaboration agreement with the IT organization at each university so that faculty get local support and the local IT staff works with UF IT Research Computing staff to address issues that go beyond what can be done by the local staff at each institution. This results in a distributed system of ownership of responsibilities and management of the resources and the support and training. UF provides the basic resources, but there are resources committed at each institution. This model is complementary to and distinct from the NSF XSEDE operation as well as the commercial cloud model. The experience during the first few months of operating this way is positive, but the period is too short for a definitive evaluation.

The comment address **C. Security requirements.**

Build the NAIRR to be compliant with security controls for restricted data from the start, such as 800-171 as required by DFARS.

Context: During the first months of operation of the HiPerGator AI, we were immediately faced with the need to handle vast amounts of restricted data for which extra security and compliance controls were required in order for the project to be allowed to move forward. Fortunately, UF has been operating a secure computing environment for some years that is compliant with NIST 800-53 moderate and NIST 800-171, which allowed UF to meet the restricted data requirements.

Question 2. Which capabilities and services (see, for example, item D below) provided through the AIRR should be prioritized?

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.

All these capabilities are required as well as training and consulting services.

Context: In the experience obtained during the first year of the UF AI Initiative, we have learned that researchers and students need all of the above including access to people knowledgeable in AI to build both introductory training as well as advanced training on AI and AI tools and consulting to teams who want to embark on a project (or write a proposal) that uses AI in any subject matter domain. The training is needed in the form of short videos, half-day workshops, and as support to faculty teaching semester-long courses involving AI. The majority of researchers and students need access to single GPUs, in which case the scaling goes with the number of users. In addition, there is the need to support very large AI problems that need the entire 140-node SuperPOD for several days to perform the desired machine learning on massive data sets.

Question 3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Governance of the NAIRR needs to include oversight from a broad community, unlike other computing resources that are well-managed by governance boards and review panels consisting of mostly technical experts.

Context: The NAIRR will have a governance structure of the NAIRR that sets strategic direction, defines the processes around resource allocation, including setting standards and requirements for proposals requesting access to the NAIRR.

The NAIRR must first determine whether there is a body of principles of ethical and responsible research and development of AI already in existence that is suitable for the needs of the community. We think that does not exist as yet. Therefore, the NAIRR could launch a task force with university researchers, industry leaders, and governmental personnel to draft a set of guiding principles that could be vetted by the community. Upon adoption, the creation and funding of a Center in Equitable AI or AI & Ethics or similar could be tasked with studying and training based on those principles. Further, it could be tasked with vetting AI products to certify that they meet the criteria set forth for ethical and responsible research in AI. In our experience at UF, the Equitable AI initiative is a step in that direction. However, an actual Center in this area could bring together people from a much broader swath of the AI community

To enhance the ability of the NAIRR to reinforce and advance the principles of ethical and responsible research and development of AI, it is imperative that the governance structure includes representatives not only from the technical computing and AI communities, but crucially involves people with background and expertise in all aspects of life and society: that means businesspeople, accountants, auditors, lawyers, philosophers to weigh in on issues of ethics and privacy, social and political scientists, psychologists in addition to the usual people who serve in these positions for computing resources. A characteristic of AI is that it directly impacts all aspects of life, unlike many other computing activities that have indirect impact, for example the computing done as part of building cars or airplanes is not visible to the people driving cars or flying in airplanes. This requirement or representation also extends to the composition of committees and panels that are asked to review proposed use of the NAIRR.

Question 4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

The COVID High Performance Computing Consortium is a perfect example of a resource structure that can address the infrastructure requirements for the NAIRR. The same NSF and DOE supercomputing resources and the expertise to build, operate, and maintain them can be leveraged to build the NAIRR. The commercial cloud providers have infrastructure and processes that are ready to contribute to the capacity and capabilities of the NAIRR. Crucially, the resources in NAIRR are not just computing machinery and data storage systems. In addition, there is a need for education, training, and consulting in AI. Numerous universities are hiring faculty with expertise in AI. The NAIRR needs to formulate a mechanism to leverage them and their expertise as a resource.

Question 5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public-Private partnerships are crucial.

Context: The partnership between UF and NVIDIA in the AI Initiative has been extremely beneficial, beyond the obvious fact that the system was donated to UF. NVIDIA has made training materials available and new materials are being developed together. The partnership

provides access to the engineers in NVIDIA so that when researchers run into difficult problems, caused by hardware or software/applications issues, they can be addressed collaboratively in the partnership to advance the mission of education and research as quickly as possible.

Question 6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The limitations are the well-known obstacles of getting started. To make the NAIRR reach underserved communities, NSF needs to provide a complementary program like the CC* program to enable campuses to build the local infrastructure, including training and consulting support staff network bandwidth, to effectively make use of the NAIRR, no matter what its technical implementation and geographical distribution ends up looking like.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

University of Illinois Chicago

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



October 1, 2021

Wendy Wigen
National Coordination Office
2415 Eisenhower Avenue
Alexandria, VA 22314

Re: *RFI Response: National AI Research Resource*

In response to the expanding interest and opportunities to develop artificial intelligence, we support the creation of the National Artificial Intelligence Research Resource Task Force. There is enormous need for infrastructure to expand access, advance collaborations, provide training, improve equity in access and development, and to prioritize standards and practices which respect the privacy and autonomy of subjects. Per the instructions, we have compiled answers and additional considerations for the team sorted by the question numbers below. We hope that this feedback will assist in the development and planning for the Task Force.

1. What options should the Task Force consider for any of roadmap elements A through I above, and why?

- *It will be critical to ensure equitable access to resources. A current limitation for many interested in AI research is access to standard resources and computing power. To disrupt consolidation of knowledge and capacity, specific attention, policy, and planning will be required to facilitate open licensing for educational use and inclusive groups who do not already have such infrastructure. We recommend the inclusion of advocacy groups such as SPARC (<https://sparcopen.org/>) with missions that focus on advancing open knowledge and that have extensive experience in development and consideration of equity by design solutions.*
- *Broad federal agency partnership and representation will be part of developing governance (B). We especially recommend the inclusion of the Institute of Museum and Library Services – which has been funding the development of practices, preservation standards and assessment of reuse for artificial intelligence projects. Academic and public libraries will continue to be outreach and educational leaders for the NAIRR and can provide access to the resources across a span of disciplines and communities. The FTC, recently funded for greater engagement with technology and privacy, will be another critical leader for developing governance and accountability.*

- *Beyond federal agencies, other governance partnerships should be considered at the state and international level (B). Artificial intelligence programs, resource use and impact will arise at both more local and global levels. Care will be needed to ensure that resources developed by NAIRR are compliant with requirements such as GDPR and other international laws governing use of AI and individuals' information.*
- *A final consideration related to governance (B(ii)) will be developing and enforcing consequences and a plan for documented and ongoing oversight for public and private sector researchers using these resources. A plan will be needed to provide accountability and transparency related to the impact of the artificial intelligence work.*

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

- *In addition to the capabilities detailed in D, another priority should be the development of a standard reporting methodology and required description for data sets and projects. Especially for datasets provided for training algorithms, standardization will be essential to begin to identify the limitations of what the AI has been trained upon, where revision and reassessment is likely to be needed, and where the ability to predict future outcomes is more likely to be inaccurate. This follows best practices in other scientific fields where standardized reporting structures for measurements have long been used to allow appropriate comparison, testing veracity and identifying limits of reproducibility and reusability. These standards should be compiled with an international community to create resources which can be used beyond national borders to meet global challenges.*

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

- *As we develop this resource, it is important to avoid the idea that artificial intelligence is a panacea for complicated problems of public and clinical health, criminal justice and overcoming historical and current discriminatory practices. AI researchers, educators, companies and others should consider and document potential downstream reuse and the harms that may result from implementing artificial intelligence solutions. There must be mechanisms and support for researchers and the NAIRR Task Force to identify where AI is not an appropriate solution.*

- Part of reinforcing these principles includes developing and enhancing foundational training. It also requires explicit consideration and documentation of harm. Ethical engagement and issues of privacy should remain consistently addressed through training programs and centered and valued in both fundamental and continuing education.
- Further, NAIRR should engage with communities whose data may be impacted in a way that conflicts with their personal and community data autonomy. Training and documentation should consider both individual and community level harm by these artificial intelligence projects. While we have seen some preliminary standards such as the CARE Principles for Indigenous Data Governance (<https://datascience.codata.org/articles/10.5334/dsj-2020-043/>), further standard and education development should recognize any disparate impacts of artificial intelligence across communities and individuals.
- In addition to improving the foundational and advanced training requirements, as we continue to presently rely on extant datasets, we must acknowledge biases in data capture. Improved and documented transparency surrounding bias in datasets is critical to disrupt perpetuation of ongoing harm.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

- Among the resources already available is extensive funded research that has identified the discrepancies between the hypotheses and actual capabilities of artificial intelligence. This research has also identified privacy and other harmful impacts. NAIRR should identify and engage with the experts who have provided decades of research and call upon their knowledge as a fundamental building block of the program. Some potential researchers include Dr. Safiya Noble, Cathy O'Neil, Dr. Casey Fiesler, and Professor David Hoffman.
- The healthcare industry is an example where algorithms have been frequently used to guide healthcare provision and to determine decisions for reimbursement purposes. This existing format can be mined both for what has served patients well but also to identify where a focus on profit and shareholder value has introduced individual and community harms which must be addressed before artificial intelligence should be adopted in other disciplines. There are similar examples of current tools showing both benefits and harms in other disciplines like criminal justice, banking, real estate and education.

5. *What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?*

- *Public-private partnerships should be a critical part of the NAIRR, but the emphasis should focus on community benefit for those whose data is captured and used. A current partnership model is healthcare, where the majority of insurance and hospitals are private. Care should be taken to identify best practices of public-private partnerships and prevent harms. Existing standards from healthcare may serve as mechanisms to start collaborative public-private partnerships.*
- *As part of the development of these partnerships, there is a need for transparency about potential biases to develop community trust, particularly with historically and currently underrepresented, underserved, and over-surveilled communities. Prior experience will have led many of these communities to mistrust private participation. A transparent and accountable process will be necessary to develop trust and collaboration.*
- *In addition to considerations about community trust, there are considerations of the goal of optimization for AI algorithms and projects. Projects are frequently optimized either for equity or for financial performance, occasionally both. As public-private partnerships are considered, the key principles of equity must be foregrounded by NAIRR to prevent a singular focus on the extraction of intellectual property in pursuit of profit.*

6. *Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?*

- *There are limitations in terms of funding, accountability, historical data accuracy and individual and community mistrust due to the past and current disparities in algorithmic impact and benefit. There are few laws that protect individuals or communities from the short- and long-term harms of artificial intelligence and enormous potential for increased redlining. There is the threat of replication and perpetuation of disparate practices based on race, gender, sexuality, income, health status, and other forms of bias. Specific laws and federal policies that improve transparency and accountability and prevent harm are likely to be needed.*
- *An additional significant challenge is limited workforce opportunities in the public sector. For example, in healthcare, there are few jobs available to focus on ethical AI healthcare development and those that are available offer fewer resources and lower salaries than the private sector. This*

restricts the number and type of projects undertaken and the speed to develop them. In the interim, for-profit endeavors proceed without oversight over the use of historically biased datasets. To address this, targeted federal funding should be made available to promote and sustain a workforce beyond private sources.

- *Artificial intelligence is not solely a consideration of the United States, and as we consider the development of standards, training, datasets and access one consideration should be addressing potential foreign influence. The United States has an opportunity to engage globally in order to ensure that AI tools are able to serve not only our citizens but our national partners in order to meet grand global challenges.*

Sincerely,

Joanna L. Groden, Ph.D.
Vice Chancellor for Research

Andrew Boyd, M.D.
Chief Research Information Officer

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Response to “Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource” (document 86 FR 39081)

National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign

The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign has been developing AI-capable resources for its own and external research for many years. Based on this experience, in this response we are providing information in response to Questions 1 and 4 jointly, in two aspects that are both related to roadmap elements D and H. Specifically, we have learned that for both technical (roadmap aspect D) and financial reasons (roadmap element H), an Artificial Intelligence Research Resource, whether providing this capability to a single institution such as the University of Illinois or to the Nation:

- 1. It is essential to make use of multiple technology types and usage models of resources, including on-premises systems, national HPC and storage resources, and commercially-provided cloud resources, and as much as possible, to offer common interfaces to the resources.*
- 2. It is equally essential to provide expertise, including technical support, alongside the computing and storage resources.*

In the remainder of this response, we provide more details on our experiences and explain how we can come to these conclusions.

Relevant NCSA experiences

1. Cloud investigations

During the last year NCSA has worked closely with many Fortune 500 companies to understand why and how some of them were moving to commercial cloud resources for their research activities, and to better understand the differences between on premises systems and cloud systems, both of which can be used for multiple purposes including AI research. During that time NCSA has run multiple proof-of-concept tests to find ways to move data science, machine learning, and normal high-performance computing

solutions into the cloud. We took existing workflows and code and shifted them to Google Cloud Platform (GCP) to evaluate both performance and cost for these corporate partners. During this year, all basic functions of these partners' research workflows were moved and tested. We learned three key lessons:

1. The cloud platform was not built to run applications originally designed for high-performance resources. Much of the software used on cloud HPC environments is the same code developed for traditional HPC systems. However, support for that software within the cloud is typically not offered as part of the cloud solution. It is typically up to the individual to make connections to the developer of the code and to the community that supports it to run in the cloud. With the in-house solutions, the in-house staff has usually done the installation and can provide basic support for the software.
2. When testing performant file systems, access to knowledgeable storage engineers is more important when running in GCP than when using local hardware solutions.
3. When moving workflows into the cloud, we found shortcomings related to the debugging and error reporting needs of developing a new system. Running controlled benchmarks led to results that didn't match what was expected and debugging those solutions was more complex and close to impossible without planning for it in advance. Cloud resources do not keep logs and related information from one session to the next, which makes comparisons more difficult.

Overall, while the advantage of the unlimited resources seemed great on paper, the knowledge and skills to implement valuable science required much more technical staff time than advertised. On the other hand, using in-house platforms to obtain research results was much easier for graduate students and researchers without all the extra technical skills required.

Additionally, a small team at NCSA experimented with cloud computing in the context of the 2.4 PB TerraFusion project (<https://digirolamo.web.illinois.edu/projects/terra-fusion/>) in 2018, by re-running a subset of the project's processing and storing the data in AWS, then scaling the results from the subset to the entire project dataset. We learned that for this project, storage costs dominated our AWS expenses, and the cloud costs substantially exceeded the costs of equivalent services provisioned in-house.

2. HAL: an NSF-funded MRI resource for AI workloads

NCSA received an NSF Major Research Instrumentation (MRI) award to develop, deploy, and operate a purpose-built computational instrument for running AI workloads. The system, named HAL, has been operational since March of 2019, enabling over 700 users to develop and train AI models. The system consists of 16 IBM Power9 AC922 servers with NVIDIA V100 GPUs, Infiniband EDR interconnect and a DDN flash array storage. Coupled with the commercial and open-source software, the system enables researchers to achieve state-of-the-art results on many AI models. While developing and operating HAL, we learned a number of valuable lessons both about technology requirements for AI applications and challenges in using this technology to obtain state-of-the-art results.

1. The compute, storage, and interconnect components all have to be well-balanced and matched to enable a high-productivity environment. For example, our initial choice for the storage solution was only able to supply data fast enough to a handful of nodes, making it impossible to productively use the entire system. Only after deploying a substantially improved storage subsystem were we able to reap benefits of the entire system.
2. The AI software stack evolves at a very high rate, with new versions of major tools being updated literally every day. As a result, significant personnel effort is needed to keep the system current with the AI tools while maintaining compatibility and performance for developers and users.
3. The AI user community is very diverse, ranging from those developing domain-specific models to those developing actual AI systems. Our experience shows that the community needs a significant amount of help to utilize the technology efficiently. This help needs to come in a variety of forms, ranging from providing user-friendly interfaces to the complex underlying hardware and software, to old-fashioned one-on-one tutoring, troubleshooting, and software tool support. This level of user support is critical to enable the use of the AI technology.

3. Extreme-scale data analysis on Blue Waters

At-scale computing of data intensive workloads such as large-scale image analysis and AI inferencing requires support for a set of elements that don't appear at the scale of desktop, research group, departmental or campus scale systems. Projects such as Digital Elevation Model (DEM) generation and tree mapping in satellite imagery projects described below are examples where NCSA's expertise in data movement, workflow software, software containerization, and HPC programming techniques enabled domain experts to achieve unprecedented results.

The DEM project (<https://bit.ly/NCSA-AI-DEM>), a collaboration of NCSA at the University of Illinois, the Polar Geospatial Center (PGC/UMN) at the University of Minnesota, the Ohio Supercomputer Center at the Ohio State University (OSC/OSU), and the National Geospatial-Intelligence Agency, generated millions of individual stereoscopic DEMs extracted from pairs of sub-meter resolution satellite imagery by applying fully automated, stereo auto-correlation techniques to overlapping pairs of high-resolution optical satellite images using the open source Surface Extraction from Triangular Irregular Network (TIN)-based Search-space Minimization (SETSM) software developed by OSC/OSU. Images were processed with SETSM using compute power on the NCSA Blue Waters supercomputer system using open-source Swift and Parsl workflow software for task management and sub-scheduling. Globus and AWS services were used to transfer petabytes of satellite imagery and resulting DEMs. Memory use and computation were optimized with assistance from staff HPC expertise and associated tools, all of which are techniques that also need be used in analyzing large-scale datasets for AI applications.

The Sub-Saharan tree mapping project (<https://bit.ly/NCSA-AI-tree-mapping>), a collaboration of Illinois/NCSA, PGC/UMN, and NASA's Goddard Space Flight Center, used machine learning algorithms and sub-meter satellite imagery to identify and measure the crown diameter of more than 1.8 billion trees across an area of more than 500,000 square miles, or 1,300,000 square kilometers. The success of the project relied heavily on NCSA's expertise with HPC software containerization and IO best-practices to exploit the power of a large-scale file system for efficient inferencing.

4. XSEDE user support and initial Delta user support

The Extreme Science and Engineering Discovery Environment (XSEDE, <https://www.xsede.org>) is an NSF-funded virtual organization that integrates and coordinates the sharing of advanced digital services, including supercomputers and high-end visualization and data analysis resources, with researchers nationally to support science. NCSA leads the XSEDE project, which is operated in collaboration with 16 other institutions¹. There is a rapidly growing number of AI-based or AI-enhanced applications already being supported on XSEDE-allocated resources.

For the upcoming NSF-funded and NCSA-hosted Delta supercomputer system (<http://www.ncsa.illinois.edu/enabling/delta>), interviews with investigators who plan to use the system show movement in the science community from purely physics-based models to adding AI components for parts of the models or the analysis of the data from

¹ <https://confluence.xsede.org/display/XT/Subaward+PIs>

the models. These researchers typically don't have AI expertise or experience, but have decided that they need to learn these skills and use these methods to do the best research.

In both cases (XSEDE and Delta), these researchers need support, and our XSEDE experience shows that this support needs to be long-term (multiple years in many cases), in part because these researchers' needs change as they progress and learn more about their science and about potential AI methods.

Additionally, our XSEDE experience has shown that different researchers request different types of hardware resources, and that these needs change over time. Today, this typically focuses on GPUs, and specifically, many GPUs per node and large memory per node and per GPU. This leads to some XSEDE systems being in high demand for AI research, while other systems are more used for other types of computational and data research.

5. C3.ai Digital Transformation Institute

The C3.ai Digital Transformation Institute (C3.ai DTI, <https://c3dti.ai/>), a research consortium that includes a support effort conducted jointly by NCSA and UC Berkeley, has accumulated specific experience supporting research on enterprise-class cloud AI infrastructure. The C3.ai DTI is entering its second award year. From a user support perspective, a few key lessons learned after the first award year going into the second are:

- The ability to leverage multiple resources, including cloud, HPC and in this case, C3.ai, is crucial for broad research team success.
- In-house technical support that can work closely with research teams and steer them toward the appropriate platforms removes barriers for those teams.
- Close cooperation with C3 (or other cloud framework providers) is needed both to leverage their in-house expertise and, critically, to translate research goals into a company culture more accustomed to considering business goals and deliverables.
- Support teams need diverse expertise, including data science/AI/machine learning as well as software engineering, computer science, HPC, system admin and solid scientific backgrounds, to be able to provide successful support.

For C3 deployments, the close working relationship between the C3.ai DTI support team and C3.ai engineers has enabled us to build technical proficiency, develop working code, and deploy infrastructure for new research teams as they come online.

The software engineering, cloud and system administration, and diverse domain expertise of our support team have all been critical in enabling the C3.ai DTI to engage research groups on

a scientific and computational level and to develop tools to facilitate quicker onboarding and prototyping on sophisticated cloud deployments. Any effort to support large scale AI research efforts will benefit from both diverse computational environments as well as a deep and broad support team expertise.

NCSA lessons learned

From these experiences, we have learned the following two lessons:

1. It is essential to make use of multiple technology types and usage models of resources, including on-premises systems, national HPC and storage resources, and commercially-provided cloud resources, and as much as possible, to offer common interfaces to the resources.

Over multiple decades, an diverse ecosystem of resources has become available to researchers, including on-premises systems (including company, university, and laboratory compute and storage resources, whether operated as clusters or clouds), national HPC and storage resources (e.g., provided by NSF, DOE), and commercially-provided cloud resources.

Each of these types of resources has positive and negative aspects for particular use cases, including the means of access and scheduling, the scale of usage, the ease of use, models of service provisioning (IaaS, FaaS, SaaS, etc.), specific resources (number and type of CPUs, GPUs, FPGAs, TPUs, etc.), data access, the cost to the end user, and the costs and business model for the resource provider, including support staff, system costs, operations, and energy usage.

While any one use case might have a resource that is best according to some criteria, there will never be a single best resource (or even model of resources) for all use cases. Therefore, a general National AI Research Resource should support and offer access to multiple models of resources, which includes resources that are now and likely will continue to be physically distributed. Similarly, data is now and will continue to be physically distributed, including both public and private data, and use cases that are data-intensive will likely be better served by computing where the data is than forcing a user to move the data to another resource. Finally, it is imperative to compare full costs when making choices about resource investment, and specifically, not to assume that commercial clouds will always be less expensive than other options, especially when coupled with the support expertise necessary to effectively be utilized.

This distributed set of resources will need to change over time in response to the evolution of both hardware and software for AI research, and elements of it must be able to support a variety of different types of work, including basic research in AI methods as well as applied research in the use of models to address both academic and commercial problems. This requires resources that include stable “production-AI” resources that can be used for method development and applied research as well as those that support the investigation of emerging novel technologies, including experimental platforms.

2. It is equally essential to provide expertise, including technical support, alongside the computing and storage resources.

There are a relatively small number of highly experienced AI users, while many potential users of AI have little understanding of AI beyond its promise, and some college graduates who have taken AI-related coursework are quite knowledgeable about AI methods, but typically have limited exposure and experience with AI tools, especially for any practical-scale applications.

These different categories of users have different needs, such as operational support for highly-experienced users, basic training and education in AI and the use of AI platforms for inexperienced users, and tool-specific training for those in the middle, particularly as tools change. These tools are frequently deployed across many platforms (edge-to-cloud-to-user), they consist of complex workflows, and are rapidly evolving. Overall, developing applications in such environments is very demanding and challenging. Therefore it is critical to have access to the technical expertise provided by dedicated personnel in order to best utilize the systems. Additionally, experienced support staff with AI expertise can help users of all types avoid common errors that could make their results incorrect or of limited utility (e.g., biased.)

In more detail, AI researchers who want to use large-scale shared resources are confronted with a myriad of options and possibilities. Depending on the goals of the researcher and the current state of their software development for research, cloud and other shared platforms can present both obstacles and opportunities. One common scenario is a team that has developed prototype software that runs on a single local computer (laptop or small workstation) and now wants to move it to a cloud. This team will typically need support in identifying, provisioning, and then deploying their software on the appropriate resources in that environment, which is typically not available from a cloud vendor but is for in-house resources and nationally-provided research HPC systems.

Support for research users is not the same as the general support that system and software vendors typically provide. Researchers work in an iterative, question-based manner, where they typically cannot provide requirements for what they need in advance, but rather, their requirements may change after each iteration of the research cycle. Therefore, the staff who support them need to be comfortable working in this same style. Over the last ten-plus years, new roles for staff have been developed who can do this, mixing understanding of researchers and their style with professional skills and understanding of underlying systems. These roles, such as research programming or research software engineering (RSE) that specialize in software development aspects, and data science that specializes in data analysis aspects, have become essential parts of the research landscape in academia, national laboratories, and industry, and have come together in the US Research Software Engineer Association (US-RSE, <https://us-rse.org>) and the Academic Data Science Alliance (ADSA, <https://academicdatascience.org>).

More complicated scenarios involve teams that lack the expertise and time to become experts in web platform management, database administration, and web development, all of which can comprise a comprehensive research platform and may become relevant as teams scale up their efforts. For smaller research teams this is untenable. While porting existing workflows onto large-scale shared resources can enable more throughput and immediate return for researchers, many teams can benefit from a more comprehensive approach in managing their research software projects without becoming experts in all of the underlying technologies.

The use of frameworks that integrate database, web, and computation can enable better end-to-end computational and analysis workflows, especially for large scale HPC-class problems, but also introduce more abstraction and somewhat steeper learning curves than basic porting of existing workflows. A curated approach for AI in large-scale shared resources that leverages a framework-like approach should offer many on-ramps for small to medium teams of domain-experts. Some will be content with a basic computation-only environment, but many teams will benefit from a more comprehensive approach that includes integrated analysis, hyperparameter tuning, cataloging of results, and so on.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

NSF AI Institutes

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Response to the NAIRR Request for Information

Submitting Organizations:

AI Institutes established in 2020:

- [NSF AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography](#)
- [NSF AI Institute for Foundations of Machine Learning](#)
- [NSF AI Institute for Student-AI Teaming](#)
- [NSF AI Institute for Molecular Discovery, Synthetic Strategy, and Manufacturing](#)
- [NSF AI Institute for Artificial Intelligence and Fundamental Interactions](#)
- [USDA-NIFA AI Institute for Next Generation Food Systems](#)
- [USDA-NIFA AI Institute for Future Agricultural Resilience, Management, and Sustainability](#)

AI Institutes established in 2021:

- [NSF AI Institute for Collaborative Assistance and Responsive Interaction for Networked Groups](#)
- [NSF AI Institute for Advances in Optimization](#)
- [NSF AI Institute for Learning-Enabled Optimization at Scale](#)
- [NSF AI Institute for Intelligent Cyberinfrastructure with Computational Learning in the Environment](#)
- [NSF AI Institute for Future Edge Networks and Distributed Intelligence](#)
- [NSF AI Institute for Edge Computing Leveraging Next-generation Networks](#)
- [NSF AI Institute for Dynamic Systems](#)
- [NSF AI Institute for Engaged Learning](#)
- NSF AI Institute for Adult Learning and Online Education
- [The USDA-NIFA Institute for Agricultural AI for Transforming Workforce and Decision Support](#)
- [The USDA-NIFA AI Institute for Resilient Agriculture](#)

Background: The National AI Institutes Program established **18 AI institutes in 2020 and 2021, with a total investment of \$360M over 5 years**. Each institute covers both foundational and applied (or “use-inspired”) AI research, and they collectively address the use of AI for a wide range of societal benefits. These institutes collectively support **over 500 researchers in over 40 states**, in a broad range of areas of core AI research, including foundational machine learning, scientific machine learning, autonomous systems, computer vision, natural language processing, multimodal interaction, edge intelligence, optimization, efficient AI, trustworthy AI, interpretable AI, fair/unbiased AI, ethics in AI, and AI education, as well as numerous broad application domains, such as agriculture, computer system design, education and adult learning, engineered systems, home healthcare, molecular manufacturing, nutrition, physics, and weather and oceanography. Over the next five years and probably much longer, the AI-focused research, education, and outreach needs of these 18 institutes will represent a very large and highly impactful user community for the NAIRR. The responses below represent the combined needs of these highly multidisciplinary AI Institutes.

Most importantly, the intensive planning and discussions that went into creating these 18 institutes have uncovered numerous important requirements for multidisciplinary AI research, some (but not all) of which are summarized briefly in the responses below. Our overarching process recommendation is to consult a representative set of individuals from these 18 institutes in the planning and creation of the NAIRR.

Q1. What options should the Task Force consider for any of roadmap elements A through I, and why?

- **(Item A) Goals:** We believe that the NAIRR should consider the needs of both foundational and multidisciplinary use-inspired AI research for a wide range of target domains. It should ensure that competitive processes for access to these resources are folded into competitive processes for research funding so that funded research teams are guaranteed access to the computational resources they need at the same time they are awarded other resources. It should consider all tiers of resources, spanning from high-end training to low-end edge devices and everything in between.

(Item A) Metrics: Metrics of success should span from the foundational to the applied. Some important metrics include (a) measures of cost, e.g., in comparison with publicly accessible commercial alternatives; (b) measures of access, spanning race/ethnicity/gender, career stages, university type; (c) measures of data set availability and usefulness, e.g., through scientifically valid surveys of various potential user communities; (d) measures of research successes, such as important research breakthroughs, standards adopted, new collaborations, etc.; (e) measures of performance benefits, e.g., on standard AI benchmarks or in quantifiable use-inspired applications; (f) measures of commercial impact, e.g., technology transfer, licensing, or startups that are spun off; and (g) qualitative results of an annual review process, conducted by a respected team of independent external evaluators.

- **(Item B) Responsible agency or organization:** We would recommend NSF as the main agency to create and govern this resource. Their mission addresses a very broad range of foundational and applied scientific research, they already fund large national computational resources, and they have previously coordinated similarly broad programs.
- **(Item C) Governance, Oversight, Resource Allocation:**
 - The governing body should include a three-way public-private partnership of academic leaders, representatives from federal funding agencies, and industry. A non-profit entity should serve as the coordinating entity, with oversight by the combined set of above stakeholders.
 - The NAIRR governing body should include subgroups with expertise in important subareas of AI as well as key application domains. The NAIRR should leverage the entire research and development community, by enabling community contributions of software tools, curated data sets, tutorials and usage manuals, discussion forums, evaluation studies, etc., all supervised by appropriate NAIRR experts.
 - Resource allocation processes should not separate research funding (e.g. NSF grants to

support PIs and other personnel in the conduct of research projects) from the computational resources needed to undertake that research. Access to the NAIRR should be allocated as part of the ordinary funding process for research programs and without requiring a separate proposal.

- For those without existing research funding or proposals in review, an open process similar to XSEDE proposals could be used, but with much more flexibility, e.g., without an expectation of predictable milestones or a timeline of outcomes, and also without requiring significant existing HPC experience, which blocks new research users.
- **(Item D) Capabilities:** We recommend a number of specific capabilities for the NAIRR:
 - **High scalability** in both compute capacity and memory size, supporting both traditional parallelism e.g., using MPI/OpenMP, and heterogeneous parallelism for machine learning, e.g., GPUs and deep learning (DL) accelerators.
 - **Configurable edge computing** ecosystems that enable experimentation with a wide range of edge AI applications, spanning very low-end systems (e.g., in IoT devices) to intermediate scale (e.g., mobile computing or AR/VR/XR) to higher-end systems (e.g., autonomous vehicles or smart cities or smart manufacturing).
 - **Data sets and data management workflows:** Curated high-quality labeled and unlabeled data sets spanning as many application domains as possible, as well as an easy-to-use pipeline for data cleaning, curation, imputation, and fusion. Outreach and awareness (along with sustained support) to identify, disseminate and grow these data sets.
 - **Benchmarking:** Exclusive and/or isolated access to systems or subsystems for benchmarking.
 - **Future-proofing:** Increasing scalability requirements over time as community needs and use cases expand.
 - **Training materials:** Extensive training materials for novices and experts alike, including training workshops, challenge competitions, and curriculum development.
 - **User support and ease of use:** Compute resources that are easy to access, understand, configure and use. Technical support for users who need assistance with getting started using the NAIRR, or with using it more efficiently and effectively.
- **(Item E) Barriers to dissemination and use:** *No response provided.*
- **(Item F) Security requirements:** CUI (Controlled Unclassified Information) data access policies, e.g., for energy-domain data, HIPAA and FERPA protected data sets, proprietary commercial and end-user data sets, secure storage and processing of proprietary or sensitive data.

- **(Item G) Anonymity, privacy, and accountability:** End-users should enjoy strong privacy and intellectual property (IP) protection guarantees for their identity, activities, results, and outcomes. Team members *and* core community contributors (selected by a well-specified public process) should be publicly identified and accountable for their contributions. NAIRR representatives should be available to answer questions about policies and practices.
- (Item H) As emphasized in the response to Q5, below, the private sector should be encouraged and even expected to play a major role in funding the creation, maintenance, and long-term sustainability of the NAIRR, and in tracking and acknowledging the commercial benefits that accrue from it.

Q2. Which capabilities and services provided through the NAIRR should be prioritized?

The most important capabilities and services the NAIRR should prioritize are:

- **Personnel funding:** AI infrastructure and data sets are not just about the "big iron": there is a large "people expense" and the NAIRR should be funded to support this component of the expense.
- **Datasets:** Dedicated funding for creating/curating/annotating/maintaining datasets is critical.
- **Training and user support:** Training -- including online, remote tutorials -- for use of the computational and data resources is essential. So is a technical support team dedicated to assisting both novice and advanced users to make best use of the resources.
- **System heterogeneity:** The computational infrastructure should include a diverse range of hardware, software, and networking, e.g., CPUs, GPUs, FPGAs, and accelerators, spanning high-end to low-end systems.
- **Scholarships:** Student funding for diversifying the user base of the NAIRR will be important.

Q3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

- **Ensure a diverse governing body:** Ensure that the group and agency governing this resource is diverse in all respects (not just race/ethnicity/gender but across university type, career stages, etc). This will bring a wealth of ideas to ensure a level playing field for all in terms of access.
- **Develop and publish an Ethics Policy:** The starting point is establishing a culture of ethics. This can be expressed in a clear policy statement that defines expectations for the ethical and responsible use of AI, enforceable for projects that make use of the NAIRR. Include requirements about the contents of data sets, with a particularly high bar for data that could become part of a formal or de facto standard or otherwise have potential to significantly impact equity and justice.

- **Use equitable and transparent allocation policies:** A traditional, peer-reviewed, proposal-driven system will tend to favor established researchers. If there is a proposal-driven process, it should include mechanisms (e.g., a “junior researcher” track or a track for non-R1 institutions) that allocate a defined fraction of the resources with substantial average allocations.
- **Provide incentives:** Explicitly reserve a fraction of the resources for research programs that focus on enhancing ethical and responsible use of AI. Other funding programs that use the NAIRR should be required to encourage these principles, and provide supplementary support for projects that successfully demonstrate that they have adhered to these goals.
- **Insulate researchers from pressure:** Excessive reliance on *direct* industry funding can undermine important ethics-related research and education in AI. Impartial, federal funding sources such as NSF are crucial. Industry contributions should be channeled through federal agencies with carefully constrained industry participation.
- **Require adherence:** The NAIRR should also require that (many or most) projects using the NAIRR include an explicit component to measure and/or train people in ethical and responsible use of AI.
- **Include a strong ombudsman:** Fund and support an independent ombudsman with authority to investigate and penalize projects that violate the Ethics Policy.
- **Be accountable:** Measure and publicize key metrics regarding the NAIRR’s user community, including equity, bias, diversity, and other important goals.

Q4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

- **Existing large-scale systems:** Two obvious examples are the high-end computing resources of the NSF and DOE-funded supercomputer centers (e.g., Delta, Cheyenne, Stampede2, Frontera) and private cloud services (Google Cloud, AWS, Microsoft Cloud). Learning from all of this is highly recommended but you should also see what has worked and not worked well, through surveys of the users. Some of these approaches do not scale well and do not provide sufficient resources to actually achieve success while others do.
- **Large testbeds for evaluation:** Large existing research testbeds can provide a strong basis for experimental evaluation of AI-driven modeling and prediction techniques in diverse application domains. Some examples include the four city-scale research testbeds for advanced wireless research funded by the \$100M PAWR project; [NIST laboratory testbeds](#) ranging from manufacturing robotics to software tools to fire protection and others; the [SoyFACE facility](#) for climate resilience in agricultural crops; the [ESnet](#) 100G SDN testbed for evaluating advanced networking research; NSSL’s [Hazardous Weather Testbed and National Weather Radar Testbed](#); and numerous others.

- **Data sets from large public testbeds:** In addition to experimental evaluation, the above testbeds may also be valuable sources of labeled and unlabeled data sets for machine learning research.
- **Commercial data sets:** Private sector companies have vast proprietary datasets and are usually unwilling to share them, but may be incentivized to share representative samples, e.g., with some NAIRR funding allocated for this purpose.
- **Centralizing coordination of existing data sets:** Many efforts, e.g., federally funded projects, generate valuable data sets and have a requirement (or a goal) to make these public, but there are few central directories or organizational structures for coordinating and supporting access to these resources (the [National Data Service](#) may be one example). One component of NAIRR funding could be allocated to publicizing the need for such data sets, and to identifying, curating, coordinating, and supporting access to related datasets from disparate sources.

Q5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

US and international companies have benefited tremendously from international research efforts in artificial intelligence, and they should be strongly encouraged to play a major role in funding the NAIRR. Many of the most promising applications of AI are in large, highly profitable, and fast-growing industry sectors that are likely to see major gains in new technological capabilities, new product categories, and large new revenue streams. *The NAIRR will play a critical role in enabling research into such applications over the next decade or more. Large, highly profitable companies today pay virtually nothing for these benefits, and that situation must change.* Moreover, part of the reason for this situation is the lack of careful accounting of the financial impact of past research, and that should be rectified as well. **Some ways in which companies could, and *should*, contribute to the NAIRR include:**

- **Pay for it:** Substantial industry contributions to joint federal+private funding programs to support the NAIRR. For example, four companies contributed \$5M each to five national AI institutes in the 2021 round, representing 20% of the total funding for these five institutes. Other NSF-funded p
- **Donate to it:** No-cost donation of storage and/or compute time on large cloud-based services and edge-compute hardware and networking for urban and rural applications, including access to both general-purpose and specialized hardware. Private entities could also contribute in-kind resources, or provide access to production systems or valuable datasets, through “grand challenge” competitions. programs, such as PAWR, RINGS, MLWins etc., have also involved public-private partnerships.

- **Create data sets for it:** Creation and publication of non-proprietary data sets that are representative of commercial workloads and assistance to researchers in accessing and using these data sets.
- **Support and diversify the students using it:** Scholarship and internship programs that expand access and skills among underserved groups.
- **Encourage entrepreneurial outcomes from it:** Venture capital to incentivize spinoff companies that emerge from NAIRR-enabled research.
- **Track and publicly acknowledge new revenue streams that accrue from it:** Companies should explicitly account for and publicly acknowledge new technologies, new product categories, new services, and new workforce training enabled by NAIRR, including a public summary of the revenues and profits that accrue from these new capabilities.

Q6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

There are roughly 3,900 accredited colleges and universities in the U.S. Democratized access requires an open infrastructure, along with federation of existing agencies' assets with this infrastructure. Some suggestions for democratizing access to the NAIRR are:

- The NAIRR should design its funding structure and competitive grant programs to spread the resources across a wide range of educational institutions, along several dimensions: research capacity (e.g., R1 research universities, 4-year colleges, and community colleges), demographics, geographic location, etc.
- One key obstacle for smaller institutions is lack of research funding. Access to NAIRR should *not* require existing research funding. On the contrary, it should provide explicit access to individuals without such funding who can use the system to bootstrap their research programs.
- Research proposals for federal funding programs that make use of the NAIRR should be required to include a strong BPC component focused on enabling training and access to the NAIRR.
- The NAIRR can provide automatic access to peer-reviewed seedling programs, like the NSF CISE Research Initiation awards, that are limited to junior researchers who have not yet secured independent federal funding.
- All publications describing research that uses the NAIRR should be made freely available to the community, without barriers of paid subscriptions.

NAIRR and RFI Links

- [NAIRRTF: THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH RESOURCE TASK FORCE](#)
- [Request for Information \(RFI\) on NAIRR Implementation Plan](#)

Questions for Response

1. What options should the Task Force consider for any of roadmap elements A through I, and why?
2. Which capabilities and services provided through the NAIRR should be prioritized?
3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?
4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?
5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?
6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Thomas Yankeelov

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.

Subject: RFI Response: National AI Research Resource
Date: Friday, September 3, 2021 5:06:28 PM

To whom it may concern,

I write to offer my thoughts on the Implementation Plan for a National Artificial Intelligence Research Resource. I am concerned that this effort will place very little emphasis on mechanism-based models, and will continue to emphasize data-driven, statistical models--a trend seen throughout federal funding programs. In particular, this is an enormous problem in the medical sciences and I think it is a cause for concern. In the biomedical world, artificial intelligence relies on properties of large groups of people that hide characteristics of the individual patient — this is especially problematic for diseases that manifest themselves so differently from person to person. Without a sound mathematical theory built on the underlying physics, chemistry, and biology, it is very difficult to understand the phenomena under investigation at a deep level--and this places fundamental limitations on how to guide interventions (for example, how to treat a cancer patient; see, for example, <https://thehill.com/opinion/healthcare/463656-what-if-we-had-a-mathematical-equation-for-cancer?rl=1>).

Thus, I would simply suggest (really, beg) that the Task Force consider including mechanism-based modeling as part of this effort.

Thank you for your time and I hope you and yours are happy, healthy, and vaccinated. :)

Sincerely,

Thomas Yankeelov, Ph.D.
W.A. "Tex" Moncrief Chair in Computational Oncology
Director, Center for Computational Oncology, Oden Institute for Computational Engineering and Sciences
Director, Cancer Imaging Research, Livestrong Cancer Institutes
co-Leader, Quantitative Oncology Research Program, Livestrong Cancer Institutes
Adjunct Professor of Imaging Physics, MD Anderson Cancer Center
Professor of Biomedical Engineering, Diagnostic Medicine, and Oncology
The University of Texas at Austin

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

**Amitha Domalpally, Roomasa Channa,
Pila Ossorio**

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



September 29th, 2021

To:

The White House Office of Science and Technology Policy and National Science Foundation

RFI Response: National AI Research Resource

From:

Amitha Domalpally, MD, PhD, Department of Ophthalmology and Visual Sciences

Roomasa Channa, MD, Department of Ophthalmology and Visual Sciences

Pilar Ossorio PhD, JD, Professor of Law and Bioethics

University of Wisconsin, Madison

As researchers active in the field of AI in ophthalmology, we are appreciative of the opportunity to provide feedback to the National Artificial Intelligence Research Resource (NAIRR) Task Force. With grass root level experience in AI research, we are enthusiastic about the potential of AI to improve health systems and patient care. We are also aware of the obstacles in applying algorithms to patient care and can bring forth diverse perspectives to the implementation roadmap. We are hopeful that the guidelines provided by this task force will have direct impact and propel AI research for future researchers in many ways. Our specific comments are listed below:

2. Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized? <https://www.federalregister.gov/d/2021-15660/p-31>

The introduction of artificial intelligence (AI) to health sciences has transformed the detection and quantification of pathology from medical imaging bringing in a new era of diagnostics. The last 3 years have seen an explosion of research comparing clinicians' diagnosis to diagnosis by AI, with AI often showing superior diagnostic performance.

The initial excitement with AI has now settled into a phase of realistic expectations as limitations in performance of AI algorithms are recognized when they are implemented into real world clinical care. Editorial comments on algorithmic bias, lack of generalizability and real-world performance are trending. It took an enormous amount of research effort, time, and money for us as researchers to realize that AI does not have to be autonomous or unsupervised in order to favorably impact patients care and we are now ready for a pragmatic approach to AI. Our recommendations for prioritization include the following

1. **Resources for Study Design:** The pathway for FDA approval of therapeutics is well established with a phased approach involving defined goals, indications of use and a critical examination of results at every stage. Apart from basic concepts such as training and testing, AI related research is relatively unstructured. In many situations, with independent researchers, training and testing is performed using a limited dataset and if fortunate, ends in a cul-de-sac of publications. There is much AI research for the sake of hashtag AI. Education in the field of AI research in terms of study design, sample size, types of validation, data fairness and bias are needed to ensure that high quality research, that can be readily translated to improve health care, is conducted.
2. **Interdisciplinary Collaboration:** Collaboration between clinicians and AI developers is important at all levels of AI model development, particularly in identifying real world challenges and addressing them from the outset. Even companies such as Google that have access to unlimited data have had major setbacks when implementing their diabetic retinopathy algorithm. A comparison of multiple deep learning models for diabetic retinopathy uncovered disparities in real world performance.
3. **Availability of Standardized Imaging Datasets:** Ophthalmology has been a high yield research area for AI applications. Autonomous diagnostic system for diabetic retinopathy is one of the first FDA approved and one of the most promising AI tools in medicine. Ophthalmology is fortunate to have many publicly available datasets of ocular images for AI training. However, many of these have incomplete metadata and issues with image quality and annotations. Ocular image format is another area that severely lacks harmonization with proprietary formats rampant with all the manufacturers. Annotation is a difficult time-consuming task and crowd sourcing is not a practical option for medical imaging. Most research groups resort to the low-quality public datasets or develop their own training datasets, with unclear labeling methods. Testing is also performed on available retrospective data from clinical care with clinician diagnosis as reference standards. Even though this process is only a few years old, it still seems like the dark ages of AI research. **Standardized, well curated training**

and testing datasets are needed to prevent further disarray in the field of AI research.

4. **Creation of a Central Repository:** Many areas of AI research have central repositories for testing model accuracy, e.g., the Stanford Question Answering Dataset (SQuAd) or ImageNet. There is no such central repository in ophthalmology which results in multiple independent models being developed. Researchers work in silos and often do not share their algorithms. In addition, there are no pretrained algorithms that can be used. Providing a central repository offers opportunities for collaboration and prevents duplication of effort.

The NAIRR task force has rightly identified that a major overhaul is needed in the AI research process, particularly in the area of universal access to curated datasets. Creation of independent datasets of images and other clinical data that are considered benchmark, with constant improvement based on user feedback is imminently needed. While many of the general guidelines in medical imaging and AI can be implemented in ophthalmology, representation of ophthalmology as a field in the implementation roadmap is required to address the distinctive imaging requirements. Health Data Research of UK has taken this step and identified ophthalmology as a designated field <https://www.insight.hdrhub.org/>.

Ophthalmology is an image intensive field with many advanced imaging techniques that allow cellular level visualization. Ocular imaging equipment are specialized and developed specifically for imaging the eye. It is also a unique field of medicine where images are not interpreted by radiologists. Clinical images are interpreted by ophthalmologists and research images acquired as part of clinical trials are interpreted by certified graders (non-ophthalmologists) in reading centers. The latter is considered reference standard for FDA approval of AI algorithms. The eye has transparent tissue and allows direct visualization of blood vessels and neuronal tissue permitting retinal images to be used for predicting systemic diseases such as Alzheimer's and cardiovascular disease. AI research in ocular imaging is positioned to be a game changer and can provide a platform to provide a proof-of-concept sandbox that can be translated to the larger body of medical research.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?
<https://www.federalregister.gov/d/2021-15660/p-32>

AI has the potential to dramatically improve diagnosis, prognosis, and image analysis. AI can leverage real-world data to complement and expand knowledge gleaned from clinical trials. However, to ensure that AI realizes its full potential in healthcare, the processes for developing and deploying it must be guided by and embody ethical norms. We define the ethical approach for AI research on three broad principles - non-maleficence, respect for persons, and justice.

Non-maleficence (do no harm): This principle is generally associated with ethical and regulatory requirements that researchers minimize risk of harm to data subjects. Research-related harms can arise from the processes of constructing and updating research datasets, or from the training and deployment of AI in healthcare. Historically, the research community has used anonymity of data subjects as the mechanism for lowering informational risks in research. Recently, there has been **extensive discussion about what counts as sufficiently anonymized or de-identified data and whether high dimensional health data (including whole genome or whole exome sequences) ever should be treated as de-identified.** The risk of re-identification increases as machine learning models become more sophisticated, computing power becomes more available, and datasets become more ubiquitous and higher dimensional. A dataset intended for widespread sharing should have utility for a variety of researchers and research projects over a fairly long time span. The variety of authorized uses of a dataset, and the unforeseen future uses, mean that both data subjects and researchers have a high degree of epistemic uncertainty concerning the potential harms or benefits that could arise from authorized uses, or from misuse.

Respect for Persons: This principle is prominent in discussions of **autonomy** and **informed consent**. Respecting persons requires that researchers promote participants' autonomy and informed consent is a means for doing so. Individual informed consent has been a mainstay of traditional clinical research. With the advent of AI, some scholars have argued that individual informed consent for research using electronic health records is impracticable (prohibitively time consuming and expensive), likely to bias results, likely to slow or hinder the development of AI tools, and unnecessary for low-risk research.

As both an ethical and regulatory matter, consent can be waived under some circumstances even though doing so may lessen the degree to which research promotes autonomy. In the U.S., researchers have never been required to seek consent or HIPAA authorization for using *de-identified* data collected in the clinical context. Despite the widespread and longstanding use of health record data for research, without individual informed consent, this practice is controversial as demonstrated by the ubiquitous negative response to the story of HeLa cells and by empirical studies. Black, Indigenous, and People of Color (BIPOC) have long been underrepresented in biomedical research, and this problem could be exacerbated because they have greater concerns about data sharing than non-BIPOC.

The 2017 revision of the Common Rule introduced a requirement for "broad consent" for the storage, maintenance, and secondary research use of *identifiable* private

information or identifiable biospecimens, including material collected in the clinical context. However, under HIPAA and the Common Rule whole-genome sequences and retina scans can be treated as de-identified even though they intrinsically individuate a person. One concern about requiring broad consent for the research use of clinical data and specimens is that health care institutions will need to track such consents for all their patients, and most institutions currently lack infrastructure and procedures for doing so.

Justice: Among other things, justice in AI research has to do with the fair distribution of the science’s burdens, potential harms, and potential benefits; fair access to research, including fair processes for choosing research questions and for determining eligibility; minimizing pernicious bias to the extent possible, in both datasets and the algorithms trained on them; and maximizing the generalizability of algorithms to the extent reasonably possible.

Race, gender, and class biases have been well-documented in AI systems, including in AI for healthcare. Bias can occur at any stage of the AI development process and can lead to differential performance of algorithms for people from different social groups. BIPOC are often under-represented in data available for designing AI models, potentially resulting in models that are less effective or even dangerous for them. However, people can also be included inappropriately, represented in the data in ways that reflect social inequality existing outside of the healthcare system or pernicious bias within the healthcare system. In either of these cases, training and validating algorithms on such data can normalize, naturalize, or obscure unjust disparities.

In healthcare, when an algorithm produces results that differ by race, class, gender, sexual orientation, or other social category, it can be difficult to know whether this difference reflects the health consequences of living in an unequal society, or whether the difference reflects pernicious bias the algorithm learned from the data. If AI is designed with a “black box” approach, it is harder to identify or mitigate pernicious bias. Performance of AI should be evaluated in different settings and in different populations to evaluate for bias. Performance should be reported in terms of predefined metrics e.g., sensitivity and specificity stratified by population characteristics such as age, gender, race/ethnicity. It would be useful to have metrics for assessing the effects of pernicious bias in patterning health-related datasets and assessing how those metrics change in response to different strategies for acquiring and cleaning data.

Data Governance is key to responsible AI research. We define data governance as “a strategy for the overall management of the availability, usability, integrity, quality, and security of data in order to ensure that the potential of the data is maximized while regulatory compliance is achieved, and ethical norms are respected and integrated. Data governance helps a dataset comply with the FAIR principles. A governance strategy must be operationalized through organizational rules and norms, and structures such as steering and advisory committees. **Data governance can draw on a responsible innovation framework, such as AREA - Anticipation, Reflection, Engagement, and Action.** Under the AREA framework, people engaged in research

should consider the future consequences of their activities, develop and integrate explicit mechanisms for reflexivity, identify and engage relevant stakeholders, and respond to stakeholder feedback and self-examination by making changes to procedures and organizations (action). We believe the AREA approach would help ensure that ethical issues are contextualized and thoroughly analyzed rather than treated as compliance issues. Data governance using an AREA approach must extend across the lifecycle of a dataset, from data collection to processing, curation, sharing, use, and finally to the end of the life cycle (deletion). There may be ethical and technical issues unique to certain stages of that lifecycle, and other issues that recur across the lifecycle.

As stakeholders with direct involvement in AI research, we appreciate the opportunity to have a voice in this remarkable initiative. For any further clarification, please reach out to us.

Amitha Domalpally, MD, PhD,

Assistant Professor, Department of Ophthalmology and Visual Sciences

Research Director, Wisconsin Reading Center

Director, A-EYE Unit

University of Wisconsin, Madison

████████████████████

Dr. Roomasa Channa, MD

Assistant Professor, Department of Ophthalmology and Visual Sciences

Co-Director, A-EYE Unit

████████████████████

Pilar Ossorio PhD, JD,

Professor of Law and Bioethics

University of Wisconsin, Madison

████████████████████