

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Responses

The White House Office of Science and Technology Policy and the National Science Foundation released a Request for Information on May 25, 2022 to request comment on the initial findings and recommendations contained in the interim report of the National Artificial Intelligence Research Resource (NAIRR) Task Force and particularly on potential approaches to implement those recommendations. The RFI was published in the Federal Register and the comment period was open from May 25, 2022, through June 30, 2022.

This document contains the 23 responses received from academia, the private sector, and civil society. In accordance with the RFI instructions, only the first 10 pages of content were considered for each response.

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

Table of Contents

ACT The App Association	1
American Psychological Association (APA)	12
Anthropic	18
Centre for the Governance of AI (GovAI)	26
Consumer Reports	37
Data Foundation	42
Dreifus, Greg and Caso, Luis Videgaray	46
Electronic Privacy Information Center (EPIC)	54
Engine	61
Hugging Face	65
IBM	68
IEEE - USA	76
Internet2	85
SeedAI	88
Shavit, Yonadav; Kaushik, Divyansh; Lipton, Zachary C.; Bowman, Samuel R.; and Goldner, Kira	92
Sheehan, Matt; Critch, Andrew; Jackson, Krystal; and Feldgoise, Jacob	98
Software & Information Industry Association (SIIA)	102
Stanford Institute for Human-Centered Artificial Intelligence (HAI)	111
The MITRE Corporation	115
U.S. Chamber of Commerce Technology Engagement Center	127
University of Arizona, CODATA Center of Excellence in Data for Society	130
University of Southern California (USC) Information Sciences Institute (ISI)	140
Wehbe, Joseph	144
Wieder, Robin	155

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

ACT | The App Association

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

June 30, 2022

Attn: Jeri Hessman
National Coordination Office for Networking and Information Technology Research and Development
Office of Science and Technology Policy
2415 Eisenhower Avenue
Alexandria, Virginia 22314

RE: Comments of ACT | The App Association to the Office of Science and Technology Policy on its Request for Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force

ACT | The App Association (App Association) appreciates the opportunity to submit views to the Office of Science and Technology Policy (OSTP) on implementing findings and recommendations of the National Artificial Intelligence Research Resource (NAIRR) Task Force, which provides guidance to federal agencies to inform the development of regulatory and non-regulatory approaches regarding technologies and industrial sectors empowered or enabled by artificial intelligence (AI) and ways for agencies to reduce barriers to the development and adoption of AI technologies.¹ The App Association generally supports the implementation of NAIRR Task Force’s findings to support and facilitate AI research and infrastructure development by prioritizing and providing sufficient funding, while ensuring adequate incentives (e.g., streamlined availability of data to developers, tax credits) are in place to encourage private and non-profit sector research. Transparency research should be a priority and involve collaboration among all affected stakeholders who must responsibly address the ethical, social, economic, and legal implications that may result from AI applications.

The App Association represents thousands of small business software application development companies and technology firms that create the technologies that drive internet of things (IoT) use cases across consumer and enterprise contexts. Today, the ecosystem the App Association represents—which we call the app economy—is valued at approximately \$1.7 trillion and is responsible for 5.9 million American jobs. Alongside the world’s rapid embrace of mobile technology, our members create the innovative solutions that power IoT across modalities and segments of the economy. The NAIRR Task Force’s findings, and the efforts of numerous agencies with respect to AI policy and regulation, directly impact the app economy. We support the Administration’s goal of ensuring the United States leads the world in technologies that are critical to our economic prosperity and national security, and to maintain the core values behind

¹ <https://www.federalregister.gov/documents/2022/05/25/2022-11223/request-for-information-rfi-on-implementing-initial-findings-and-recommendations-of-the-national>

America's scientific leadership, including openness, transparency, honesty, equity, fair competition, objectivity, and democratic values.²

AI is an evolving constellation of technologies that enable computers to simulate elements of human thinking – learning and reasoning among them. An encompassing term, AI entails a range of approaches and technologies, such as Machine Learning (ML) and deep learning, where an algorithm based on the way neurons and synapses in the brain change due to exposure to new inputs, allowing independent or assisted decision making. AI-driven algorithmic decision tools and predictive analytics are having, and will continue to have, substantial direct and indirect effects on Americans. Some forms of AI are already in use to improve American consumers' lives today – for example, AI is used to detect financial and identity theft and to protect the communications networks upon which Americans rely against cybersecurity threats.

Moving forward, across use cases and sectors, AI has incredible potential to improve American consumers' lives through faster and better-informed decision making, enabled by cutting-edge distributed cloud computing. As an example, healthcare treatments and patient outcomes stand poised to improve disease prevention and conditions, as well as efficiently and effectively treat diseases through automated analysis of x-rays and other medical imaging. AI will also play an essential role in self-driving vehicles and could drastically reduce roadway deaths and injuries. From a governance perspective, AI solutions will derive greater insights from infrastructure and support efficient budgeting decisions. An estimate states AI technological breakthroughs will represent a \$126 billion market by 2025.³

Today, Americans encounter AI in their lives incrementally through the improvements they have seen in computer-based services they use, typically in the form of streamlined processes, image analysis, and voice recognition (we urge consideration of these forms of AI as “narrow” AI). The App Association notes that this narrow AI already provides great societal benefit. For example, AI-driven software products and services revolutionized the ability of countless Americans with disabilities to achieve experiences in their lives far closer to the experiences of those without disabilities.

Nonetheless, AI also has the potential to raise a variety of unique considerations for policymakers. The App Association appreciates the efforts to develop a policy approach to AI that will bring its benefits to all, balanced with necessary safeguards to protect consumers.

² *Id.*

³ McKinsey Global Institute, *Artificial Intelligence: The Next Digital Frontier?* (June 2017), available at <https://www.mckinsey.com/~media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx>.

To assist the Administration, the App Association offers a comprehensive set of AI policy principles below for consideration that we strongly encourage alignment in the implementation of the initial findings and recommendations:

1. **AI Strategy:** Many of the policy issues raised below involve significant work and changes that will impact a range of stakeholders. The cultural, workforce training and education, data access, and technology-related changes associated with AI will require strong guidance and coordination. An AI strategy incorporating guidance on the issues below will be vital to achieving the promise that AI offers to consumers and our economies. We believe it is critical to take this opportunity to encourage civil society organizations and private sector stakeholders to begin similar work. The NAIRR Task Force's findings should remain a key part of the United States' overall strategy for global leadership in this critical area of technology.
2. **Research:** The implementation of NAIRR Task Force's report should support and facilitate research and development of AI by prioritizing and providing sufficient funding while also ensuring adequate incentives (e.g., streamlined availability of data to developers, tax credits) are in place to encourage private and non-profit sector research. Transparency research should be a priority and involve collaboration among all affected stakeholders who must responsibly address the ethical, social, economic, and legal implications that may result from AI applications.
3. **Quality Assurance and Oversight:** The implementation of NAIRR Task Force's findings and the U.S. approach to AI generally, should advance risk-based approaches to ensure that the use of AI aligns with the recognized standards of safety, efficacy, and equity. Providers, technology developers and vendors, and other stakeholders all benefit from understanding the distribution of risk and liability in building, testing, and using AI tools. Policy frameworks addressing liability should ensure the appropriate distribution and mitigation of risk and liability. Specifically, those in the value chain with the ability to minimize risks based on their knowledge and ability to mitigate should have appropriate incentives to do so. Some recommended guidelines include:
 - Ensuring AI is safe, efficacious, and equitable.
 - Supporting the creation of algorithms, datasets, and decisions that are auditable.
 - Encouraging AI developers to consistently utilize rigorous procedures and enabling them to document their methods and results.
 - Requiring those developing, offering, or testing AI systems to provide truthful and easy-to-understand representations regarding intended use and risks that would be reasonably understood by those intended, as well as expected, to use the AI solution.
 - Ensuring that adverse events are timely reported to relevant oversight bodies for appropriate investigation and action.

4. **Thoughtful Design:** The implementation of NAIRR Task Force’s report, and the U.S. approach to AI generally, should strongly encourage the design of AI systems that are informed by real-world workflows, human-centered design and usability principles, and end-user needs. AI systems solutions should facilitate a transition to changes in the delivery of goods and services that benefit consumers and businesses. The design, development, and success of AI should leverage collaboration and dialogue among users, AI technology developers, and other stakeholders to have all perspectives reflected in AI solutions.
5. **Access and Affordability:** The implementation of NAIRR Task Force’s report, and the U.S. approach to AI generally, should ensure AI systems are accessible and affordable. Significant resources may be required to scale systems and policymakers should take steps to remedy the uneven distribution of resources and access. Policies must be put in place that incent investment in building infrastructure, preparing personnel and training, as well as developing, validating, and maintaining AI systems with an eye toward ensuring value.
6. **Ethics:** AI will only succeed if it is used ethically. It will be critical to promote many of the existing and emerging ethical norms for broader adherence by AI technologists, innovators, computer scientists, and those who use such systems. The implementation of NAIRR Task Force’s report, and the U.S. approach to AI generally, should:
 - Ensure that AI solutions align with all relevant ethical obligations, from design to development to use.
 - Encourage the development of new ethical guidelines to address emerging issues with the use of AI, as needed.
 - Maintain consistency with international conventions on human rights.
 - Ensure that AI is inclusive such that AI solutions beneficial to consumers are developed across socioeconomic, age, gender, geographic origin, and other groupings.
 - Reflect that AI tools may reveal extremely sensitive and private information about a user and ensure that laws protect such information from being used to discriminate against certain consumers.
7. **Modernized Privacy and Security Frameworks:** While the types of data items analyzed by AI and other technologies are not new, this analysis will provide greater potential utility of those data items to other individuals, entities, and machines. Thus, there are many new uses for, and ways to analyze, the collected data. This raises privacy issues and questions surrounding consent to use data in a particular way (e.g., research, commercial product/service development). It also offers the potential for more powerful and granular access controls for consumers. Accordingly, the implementation of NAIRR Task Force’s report, and the U.S. approach to AI generally, should address the topics of privacy, consent, and modern technological capabilities as a part of the policy development process. Risk management policy frameworks must be scalable and assure that an individual’s data is properly protected, while also allowing the

flow of information and responsible evolution of AI. This information is necessary to provide and promote high-quality AI applications. Finally, with proper protections in place, policy frameworks should also promote data access, including open access to appropriate machine-readable public data, development of a culture of securely sharing data with external partners, and explicit communication of allowable use with periodic review of informed consent.

8. **Collaboration and Interoperability:** The implementation of NAIRR Task Force’s report, and the U.S. approach to AI generally, should enable eased data access and use through creating a culture of cooperation, trust, and openness among policymakers, AI technology developers and users, and the public.
9. **Bias:** The bias inherent in all data, as well as errors, will remain one of the more pressing issues with AI systems that utilize machine learning techniques in particular. Addressing data provenance and bias issues is a must in developing and using AI solutions. The implementation of NAIRR Task Force’s report, and the U.S. approach to AI generally, should:
 - Require the identification, disclosure, and mitigation of bias while encouraging access to databases and promoting inclusion and diversity.
 - Ensure that data bias does not cause harm to users or consumers.
10. **Education:** The implementation of NAIRR Task Force’s report, and the U.S. approach to AI generally, should support education for the advancement of AI, promote examples that demonstrate the success of AI, and encourage stakeholder engagements to keep frameworks responsive to emerging opportunities and challenges.
 - Consumers should be educated as to the use of AI in the service they are using.
 - Academic education should include curriculum that will advance the understanding of and ability to use AI solutions.

The policy issues raised by the NAIRR Task Force involve significant work and changes that will impact a range of stakeholders. The cultural, workforce training and education, data access, and technology-related changes associated with AI will require strong guidance and coordination across U.S. federal agencies. The App Association supports the development of national AI strategies for federal agencies, which will be vital to achieving the promise that AI offers to consumers and entire economies.

Noting our general support for the current findings and recommendations of NAIRR Task Force’s report, we offer the following suggested revisions:

- **Alignment with Other Leading Federal Policies for AI:** The implementation of NAIRR Task Force’s report should align with other federal efforts to develop AI policy, such as the National Institute of Standards and Technology’s (NIST) Artificial Intelligence Risk Management Framework, a policy being developed in close collaboration with the private sector, academia, and others for voluntary use with the goal of improving the ability to incorporate trustworthiness

considerations into the design, development, use, and evaluation of AI products, services, and systems.⁴

- **Require Agencies to Advance Thoughtful Design Principles Across AI Use Cases:** The implementation of NAI RR Task Force’s report should require design of AI systems informed by real-world workflows, human-centered design and usability principles, and end-user needs. AI systems solutions should facilitate a transition to changes in the delivery of goods and services that benefit consumers and businesses. The design, development, and success of AI should leverage collaboration and dialogue among users, AI technology developers, and other stakeholders in order to have all perspectives reflected in AI solutions. As this concept must run across sectors and AI use cases, the NAI RR Task Force should continue to incorporate guidance for agencies to advance thoughtful design principles through their approaches and actions related to AI.
- **Require Agencies to Advance Ethics in AI’s Development and Use:** The success of AI depends on ethical use. An agency’s approach will need to promote many of the existing and emerging ethical norms for broader adherence by AI technologists, innovators, computer scientists, and those who use such systems. The implementation of NAI RR Task Force’s report should:
 - Ensure that AI solutions align with all relevant ethical obligations, from design to development to use.
 - Encourage the development of new ethical guidelines to address emerging issues with the use of AI, as needed.
 - Maintain consistency with international conventions on human rights.
 - Ensure that AI is inclusive such that AI solutions beneficial to consumers develop across socioeconomic, age, gender, geographic origin, and other groupings.
 - Reflect that AI tools may reveal extremely sensitive and private information about a user and ensure that laws protect such information from being used to discriminate against certain consumers
- **Augment the Requirement on Federal Agencies for Disclosure and Transparency:** The Administration should consider further prioritizing disclosure and trust priorities when implementing NAI RR Task Force’s findings. Providers, technology developers, and vendors, and other stakeholders will all benefit from understanding the distribution of risk and liability in building, testing, and using AI tools. The implementation of NAI RR Task Force’s report should therefore clearly address liability so as to ensure the appropriate distribution and mitigation of risk and liability (i.e., those in the value chain with the ability to minimize risks based on their knowledge and ability to mitigate should have appropriate incentives to

⁴ <https://www.nist.gov/itl/ai-risk-management-framework>.

do so). Further, the NAIRR Task Force should clearly require that AI policies prioritize that those developing, offering, or testing AI systems provide truthful and easy to understand representations regarding intended use and risks that would be reasonably understood by those intended, as well as expected, to use the AI solution.

- **Support the Development of, and Access to, Open Standards Needed to Drive U.S. Leadership in AI:** The implementation of NAIRR Task Force's report should support the developer and use of voluntary consensus standards that concern AI application. The App Association strongly encourages updating the NAIRR Task Force's plan to support public-private collaboration on AI through standardization by encouraging key U.S.-based standard-setting organizations (SSOs) such as IEEE to grow and thrive. The U.S. government can support such organizations through pro-innovation policies that encourage private sector research and development of AI innovations and the development of related standards.

It is critical that the United States should ensure that such standards are accessible to innovators by promoting a balanced approach to standard-essential patent (SEP) licensing. AI technical standards, built on contributions through an open and consensus-based process, bring immense value to consumers by promoting interoperability while enabling healthy competition between innovators; and often include patented technology. When an innovator gives its patented technology to a standard, this can represent a clear path to reward in the form of royalties from a market that likely would not have existed without the standard being widely adopted. To balance this potential with the need for access to the patents that underlie the standard, many SSOs require holders of patents on standardized technologies to license their patents on fair, reasonable, and non-discriminatory (FRAND) terms. FRAND commitments prevent the owners of patents used to implement the standard from exploiting the unearned market power that they otherwise would gain due to the broad adoption of a standard. Once patented technologies incorporate into standards, it compels manufacturers to use them to maintain product compatibility. In exchange for making a voluntary FRAND commitment with an SSO, SEP holders gain the ability to obtain reasonable royalties from numerous standard implementers that might not have existed absent the standard. Without the constraint of a FRAND commitment, SEP holders would have the same power as a monopolist that faces no competition.

Unfortunately, several owners of FRAND-committed SEPs are flagrantly abusing their unique position by reneging on those promises with unfair, unreasonable, or discriminatory licensing practices. These practices, under close examination by antitrust and other regulators in many jurisdictions, not only threaten healthy competition and unbalance the standards system but also impact the viability of

new markets such as AI. This amplifies the negative impacts on small businesses because they can neither afford years of litigation to fight for reasonable royalties nor risk facing an injunction if they refuse a license that is not FRAND compliant.

Patent policies developed by SSOs today will directly impact the way we work, live, and play for decades to come. SSOs vary widely in terms of their memberships, the industries, and products they cover, and the procedures for establishing standards. In part due to the convergence associated with the rise of IoT, each SSO will need the ability to tailor its intellectual property policy for its particular requirements and membership. The App Association believes that some variation in patent policies among SSOs is necessary and that the U.S. government should not prescribe detailed requirements that all SSOs must implement. At the same time, however, as evidenced by the judicial cases and regulatory guidance, basic principles underlie the FRAND commitment and serve to ensure that standard setting is pro-competitive, and the terms of SEP licenses are in fact reasonable. Ideally, an SSO's intellectual property rights policy that requires SEP owners to make a FRAND commitment would include the following principles that prevent patent "hold up" and anti-competitive conduct:

- **Fair and Reasonable to All** – A holder of a SEP subject to a FRAND license such SEP on fair, reasonable, and nondiscriminatory terms to all companies, organizations, and individuals who implement or wish to implement the standard.
- **Injunctions Available Only in Limited Circumstances** – SEP holders should not seek injunctions and other exclusionary remedies nor allowed these remedies except in limited circumstances. The implementer or licensee is always entitled to assert claims and defenses.
- **FRAND Promise Extends if Transferred** – If there is a transfer of a FRAND-encumbered SEP, the FRAND commitments follow the SEP in that and all subsequent transfers.
- **No Forced Licensing** – While some licensees may wish to get broader patent holder should not require implementers to take or grant licenses to a FRAND-encumbered SEP that is invalid, unenforceable, or not infringed, or a patent that is not essential to the standard.
- **FRAND Royalties** – A reasonable rate for a valid, infringed, and enforceable FRAND-encumbered SEP should be based on several factors, including the value of the actual patented invention apart from its inclusion in the standard, and cannot be assessed in a vacuum that ignores the portion in which the SEP is substantially practiced or royalty rates from other SEPs required to implement the standard.

We also note that several SSO intellectual property rights policies require SSO participants to disclose patents or patent applications that are or may be essential to a standard under development. Reasonable disclosure policies can help SSO participants evaluate whether technologies considered for standardization are covered by patents. Disclosure policies should not, however, require participants to search their patent portfolios as such requirements can be overly burdensome and expensive, effectively deterring participation in an SSO. In addition, FRAND policies that do not necessarily require disclosure, but specify requirements for licensing commitments for contributed technology, can accomplish many, if not all, of the purposes of disclosure requirements.

The U.S. Department of Justice (DOJ) already encouraged SSOs to define FRAND more clearly. For example, DOJ's former assistant attorney general Christine Varney explained that "clearer rules will allow for more informed participation and will enable participants to make more knowledgeable decisions regarding implementation of the standard. Clarity alone does not eliminate the possibility of hold-up...but it is a step in the right direction."⁵ As another example, Renata Hesse, a previous head of the DOJ's Antitrust Division, provided important suggestions for SSOs to guard against SEP abuses that included at least three of the aforementioned principles.⁶ The implementation of NAIIR Task Force's recommendations should be updated to advance open standards, consistent with OMB-A119 ("Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities"),⁷ open standards and access to open standards with respect to SEPs.

The App Association appreciates OSTP's consideration of the above views. We urge OSTP to contact the undersigned with any questions or ways that we can assist moving forward.

⁵ Christine A. Varney, Assistant Att'y Gen., Antitrust Div., U.S. Dep't of Justice, Promoting Innovation Through Patent and Antitrust Law and Policy, Remarks as Prepared for the Joint Workshop of the U.S. Patent and Trademark Office, the Federal Trade Comm'n, and the Dep't of Justice on the Intersection of Patent Policy and Competition Policy: Implications for Promoting Innovation 8 (May 26, 2010), *available at* <http://www.atrnet.gov/subdocs/2010/260101.htm>.

⁶ Renata Hess, Deputy Assistant Attorney General, *Six 'Small' Proposals for SSOs Before Lunch*, Prepared for the ITU-T Patent Roundtable (October 10, 2012), *available at* <https://www.justice.gov/atr/speech/six-smallproposals-ssos-lunch>.

⁷ https://www.nist.gov/system/files/revise_d_circular_a-119_as_of_01-22-2016.pdf.

Sincerely,

Brian Scarpelli
Senior Global Policy Counsel

Leanna Wade
Public Policy Associate

ACT | The App Association
1401 K St NW (Ste 501)
Washington, DC 20005
202-331-2130

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

American Psychological Association (APA)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



AMERICAN
PSYCHOLOGICAL
ASSOCIATION

June 30, 2022

Jeri Hessman

National Coordination

Office for Networking and Information Technology Research and Development

2415 Eisenhower Avenue

Alexandria, VA 22314, USA

Submitted electronically via Regulations.gov

RE: RFI Response: National AI Research Resource Interim Report

Dear Ms. Hessman –

The American Psychological Association (APA) appreciates the opportunity to comment on the National Artificial Intelligence Research and Development Strategic Plan require for information. This request represents a step in the right direction towards ensuring that stakeholders across disciplines are represented in future efforts to deploy artificial intelligence. In addition to the comments below, APA endorses the comment submitted from the Society for Industrial and Organizational Psychology.

APA is the largest scientific and professional organization representing psychology in the U.S., numbering over 133,000 researchers, educators, clinicians, consultants, and students. For decades, psychologists have played a vital role in the development and deployment of technologies and neurological science. These contributions have been essential to the currently available artificial intelligence enabled technologies and psychological science should continue to be at the heart of strategic planning of AI deployment.

The comments below represent three primary areas where the *Strategic Plan* should ensure the discipline of psychology is included: increased investments in research on artificial intelligence, ethics of artificial intelligence, and artificial intelligence and implicit bias.

750 First Street, NE
Washington, DC 20002-4242
(202) 336-5800
(202) 336-6123 TDD

Web: www.apa.org



AMERICAN PSYCHOLOGICAL ASSOCIATION

- a. Vision for the NAIRR. Including strategic goals and objectives, composition, and user base. (Chapter 2 of the report)
 - APA applauds the focus of NAIRR on civil rights, privacy and civil liberties. To achieve this goal, we encourage the NAIRR to create mechanisms within their workflow that provides frequent opportunities for stakeholder feedback and engagement. These engagement opportunities will provide organizations like the APA the ability to gather feedback from our diverse set of members to share with the NAIRR.
 - APA strongly supports the need for additional investments in research related to Artificial Intelligence. From the current technological and research standpoint, it is almost impossible to predict the impact of future AI-informed technologies. There is an imperative that as the technologies grow in their capabilities and prevalence, that research surrounding their impact also increases. Future research funding in this area should ensure that psychological and behavioral science is adequately represented. The impact of AI on mental and behavioral health must continue to be examined to ensure we mitigate any harmful impacts caused by new systems.
 - The NAIRR should also consider developing partnerships and channels of outreach to ensure adequate dissemination of research, funding opportunities, and findings. Channels available to organizations like the APA can be essential to ensuring the reach of the NAIRR is maximized.
 - The NAIRR must also ensure that psychology is incorporated into any review of human factors. Without the benefits of psychology science, the impact of human factors can't be robustly known.
- b. Establishment and sustainment of the NAIRR. Including agency roles, resource ownership and administration, governance and oversight, resource allocation and sustainment, and performance indicators and metrics. (Chapter 3 of the report)
 - One mechanism for performance evaluation missing from the proposal put forward by NAIRR is its achievements relative to goals of equity, diversity, and inclusion. As with other proposed performance metrics, having an outside organization periodically evaluate the NAIRR for their performance achieving stated goals of including diverse research, considering diverse populations when making recommendations, and ensuring communication is accessible by all audiences is essential to the NAIRR's sustained success.



AMERICAN PSYCHOLOGICAL ASSOCIATION

- c. NAIRR resource elements and capabilities. Including data, government datasets, compute resources, testbeds, user interface, and educational tools and services. (Chapter 4 of the report)
- d. System security and user access controls. (Chapter 5 of the report)
- e. Privacy, civil rights, and civil liberties requirements. (Chapter 6 of the report)
 - There are some fundamental research opportunities the NAIRR must investigate. AI Ethics and Psychology is an evolving discipline essential to the study of how AI learns from society and humans and how AI makes consequential decisions in critical settings.¹ Studies have demonstrated that AI automatically learns implicit biases from language corpora and accordingly perceives the world in a biased manner.² These implicit biases that have been documented in social psychology for decades include racial, gender, sexuality, ability, and age attitudes.³ 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.⁴ Alternatively, these AI methods could examine and analyze current and historical social and human cognition.⁵

¹ 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).

² 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.

³ 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>; 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>.

⁴ 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>

⁵ 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>



AMERICAN PSYCHOLOGICAL ASSOCIATION

This research program would allow for understanding how AI is co-evolving with humanity, as AI is shaping society and impacting individuals' lives in an accelerated manner and at an unprecedented scale.

- Given evidence that AI can reproduce discrimination and bias against individuals and groups, it is imperative the NAIRR 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. The early evidence shows that decisions about AI and how it is implemented reflect the world view and values of the human beings who design them and set policy for how it is used. 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.

Without the help of psychological science, we risk harming already disadvantaged populations and creating systems that perpetuate harmful stereotypes and bias. AI systems are often trained using large data sets of human attributes or demographics that have the potential to integrate biases related to gender identity, race, and other characteristics. These systems then spread the biases in their interactions with humans or other technology-informed systems, 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. Findings from this research should be incorporated into future deployments of artificial intelligence tools, especially when being funded or used by the federal government.

- f. Ideas for developing a roadmap to establish and build out the NAIRR in a phased approach, and appropriate milestones for implementing the NAIRR. Including data sets, use cases, and capabilities that should be prioritized in the early stages of establishment of the resource.

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



AMERICAN
PSYCHOLOGICAL
ASSOCIATION

g. Other areas relevant to the development of the NAIRR implementation plan.

While we remain broadly supportive of the strategic aims set forward by the *Envisioning a National Artificial Intelligence Research Resource (NAIRR): Preliminary Findings and Recommendations*, it is important that psychological and behavioral science is included in each strategy to ensure comprehensive consideration of the broad impact of AI technologies.

APA again thanks you for the opportunity to comment on this policy. If APA can be of any further assistance, please contact Corbin Evans, Senior Director of Congressional and Federal Relations, at CEvans@APA.org.

Katherine B. McGuire, MS
Chief Advocacy Officer, APA

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Anthropic

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

ANTHROPIC

June 29, 2022

Submitted electronically via NAIRR-responses@nitrd.gov

Reference: 87 FR 31914, Document Number 2022-11223

Subject: Anthropic Comment regarding “Request for Information (RFI) on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force”

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 initial findings and recommendations contained in the Interim Report of the National Artificial Intelligence Research Resource (NAIRR) Task Force.

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. Since our founding in early 2021, we’ve raised \$700 million, primarily to fund the buildout of world class infrastructure for running large-scale AI training jobs. For Anthropic, infrastructure is fundamental to the success of our organization and we invest significant amounts of capital and headcount into developing and maintaining it.

As articulated in our previous submission¹, we believe progress in AI safety and technical advancements depends on broad public participation in AI research. Unfortunately over the past several years, frontier AI research and development (R&D) has become heavily concentrated within a small number of corporate entities. The NAIRR represents an opportunity to democratize the resources required to build advanced AI systems, allowing non-commercial actors to actively participate in the R&D ecosystem and build systems that more accurately reflect the goals of society at-large.

We appreciate the opportunity to share our perspective on the Interim Report, specifically the proposed scope of the NAIRR, its resource requirements, and organizational design. At a high level, we recommend the NAIRR Task Force consider the following suggestions as it refines its plans for the Final Report:

¹ Anthropic. (2021, October). *Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses*. ai.gov. <https://www.ai.gov/rfi/2021/86-FR-39081/Anthropic-NAIRR-RFI-2021.pdf>

- Putting forward a more ambitious, detailed technical proposal for the scope of the NAIRR, in order to seize this opportunity to deeply invest in the United States’ AI research capacity **(Topic A & Topic B)**
- Distributing the majority of compute resources to a broad range of academic researchers, meanwhile allocating a non-trivial portion for research endeavors that rival industry-scale development efforts **(Topic B & Topic C)**
- Using the NAIRR not only as a repository of existing testbeds, but also as an opportunity to further increase the capacity to measure, assess, benchmark, and monitor AI systems **(Topic C)**
- Triaging NAIRR capabilities so that implementation can happen in a staged way – start with the minimum viable product (“MVP”) version of a NAIRR and iterate from there, while continuously testing for usability and efficiency of NAIRR infrastructure **(Topic B)**

A Bolder Vision for the NAIRR (Topic A) & A Concrete Technical Proposal (Topic C)

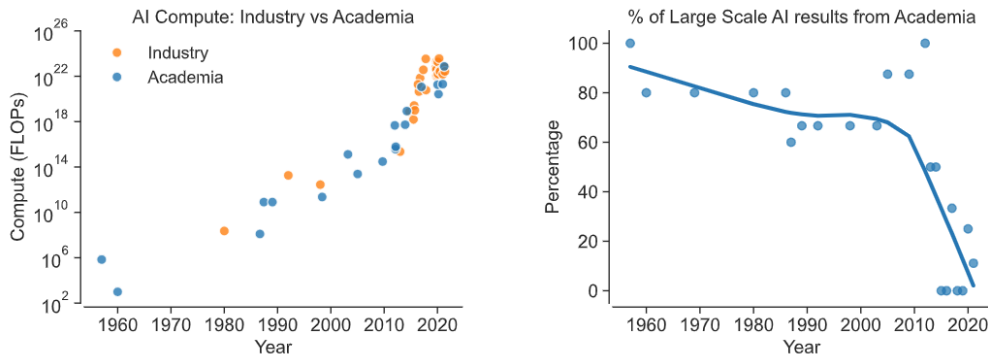
We echo the Task Force’s findings that potential breakthroughs in sustainability, national security, and other societal challenges depend on expanded access to AI resources, including compute hardware and sufficient data. The NAIRR represents an enormous opportunity for the U.S. Government to continue its longstanding support of academic research and advance AI development in the United States. While the Interim Report accurately notes the growing divide between academia and the private sector in resource accessibility (and as a result, opportunity to develop advanced systems), the compute divide is even more pronounced than what is alluded to in the report.

Compute-intensive research – the kind required to build more general-purpose, natural language and computer code processing models – is incredibly costly, and far exceeds the standard academic budget. For example, the training cost of OpenAI’s GPT-3 was estimated to be several million dollars², while Google’s PaLM was estimated to cost between \$9 million and \$23 million dollars³. As these models continue to grow in size and capabilities, the compute required to train them also increases and most academic institutions can’t allocate funding for resources at this scale. As a result, we’ve seen the industry contributions to large-scale AI research dwarf those of the academic community over the past decade⁴.

² Lambda. (2020, June 3). *OpenAI’s GPT-3 Language Model: A Technical Overview*. LambdaLabs. <https://lambdalabs.com/blog/demystifying-gpt-3/>

³ Heim, Lennart. (2022, April 5). *blog.heim.xyz*. Estimating PaLM’s training cost. <https://blog.heim.xyz/palm-training-cost/>

⁴ Ganguli, D., et al. (2022). Predictability and Surprise in Large Generative Models. arXiv. <https://arxiv.org/abs/2202.07785>



(Left) The amount of compute required by major AI projects over time is increasing exponentially for both academic (blue) and industrial (orange) projects. **(Right)** The proportion of large-scale AI results from academia is steadily decreasing. Blue curve represents a Lowess fit to the data⁴.

To meet the ambitious goal of the NAIRR to “strengthen and democratize the U.S. AI innovation ecosystem,” the U.S. Government must meet the moment with an equally ambitious and detailed infrastructure proposal. **Recommendation 3-1 of the Interim Report suggests multiple Federal agencies fund the NAIRR cooperatively, but does not point to specific agencies or detailed funding requirements.** Leaving this fundamental aspect of the NAIRR undefined leaves its potential success up to chance. We understand that funding recommendations may not be in scope for the Initial Report, but we believe it is a critical aspect to plan for in advance. We have included some estimated figures for consideration, and we urge the Task Force to include concrete recommendations and implementation plans in the Final Report. Ideally, these recommendations would be supplemented with draft legislation, which could be included as an appendix to the Final Report.

To assist in this effort, the Task Force may look to similar models implemented in Australia (National Computational Infrastructure⁵) and Canada (Advanced Research Computing (ARC) Platform⁶) to benchmark its own recommendations against the technical investments made in peer countries. However, the scope of these investments should be treated as a baseline, as year-over-year researcher demand continues to exceed available supply. In the case of Canada’s ARC Platform, only 24% of the total GPUs (a fundamental hardware component to today’s large-scale AI systems) requested were awarded to researchers in 2022 due to resource limitations⁷. If the United States wishes to sustain its leadership in AI research, it must invest in ways to make AI research more accessible. Outlining an ambitious and detailed technical proposal in the Task Force’s Final Report can be a step towards that effort.

⁵ NCI Australia. HPC Systems. <https://nci.org.au/our-systems/hpc-systems>

⁶ Digital Research Alliance of Canada. (2022). *Advanced Research Computing*. <https://alliancecan.ca/en/services/advanced-research-computing>

⁷ Digital Research Alliance of Canada. (2022). *2022 Resource Allocations Competition Results*. <https://alliancecan.ca/en/services/advanced-research-computing/research-portal/resource-allocation-competitions/2022-resource-allocations-competition-results>

Our vision of a truly competitive NAIRR for the United States would be on the order of a 100,000 GPU cluster, with an estimated cost of roughly \$4 billion for three years. This would require a significant financial investment and would represent an exemplary resource built over the course of several years. In our previous comment we contextualized this cost with the \$97 billion spent on capital expenditures⁸ by the leading U.S. digital infrastructure providers (Amazon, Google, Microsoft) in 2020. Private sector investment continues to grow: in 2021, those companies spent over \$124 billion on capital expenditures, representing an increase of 28% over the previous year⁹, and further widening the divide between resources available in industry and academia.

To align resource needs with an eventual appropriations recommendation, the Task Force may consider polling academic research departments across the United States to better understand current resource constraints and future research ambitions. Alternatively, the Task Force could partner with a think tank or non-profit organization focused on the United States' R&D ecosystem to carry out the polling exercise. In either case, we recommend the Task Force include a concrete technical proposal, informed by the needs of researchers, in the Final Report. Doing so will not only reveal the AI research potential of the academic community, but inform concrete steps towards building a stronger AI innovation ecosystem.

Create Opportunities for Large-Scale AI Experimentation (Topic B & Topic C)

Without public intervention along the lines of the NAIRR, access to compute resources will continue to be a research barrier between academia and industry. We agree with Recommendation 3-11 that resource allocation processes should be as inclusive as feasible, and **recommend the Task Force consider an additional funding and access tier for a small number of compute-intensive research projects.** As currently envisioned, the NAIRR will enable new AI research from a broader set of stakeholders, but the compute gap between academia and industry will continue to widen without a substantial investment in a handful of large-scale projects.

While the vast majority of resources would be made widely accessible to a broad range of researchers, something on the order of 30% of the NAIRR's computational capacity could be awarded to research efforts seeking to build industry-scale AI systems. Given the fact that compute-intensive models are developed almost exclusively within industry, and further investigation and access to those models is tightly controlled by corporate actors, **providing an opportunity for academic researchers to build and investigate equivalent systems directly**

⁸ Note that this CapEx figure includes spend on land, corporate offices, warehouses, etc., in addition to cloud and data center infrastructure.

⁹ Fitzgerald, C. (2022, February 16). *Follow the CAPEX: Cloud Table Stakes 2021 Retrospective*. Platformonomics. <https://www.platformonomics.com/2022/02/follow-the-capex-cloud-table-stakes-2021-retrospective/>

supports Recommendation 3-13 (“...resource allocation should be designed to incentivize contributions to the NAIRR user community or to the public good”).

We support the Task Force’s recommendation that all users of NAIRR computational resources pass a research proposal evaluation process (Recommendation 3-10), and **recommend that compute-intensive projects undergo an enhanced review that explicitly considers the potential societal impacts of such research**. To implement Recommendation 3-13 on a practical level, the proposal review process might also take into account the researchers’ intent to contribute such systems back to the NAIRR ecosystem, thereby enabling further investigation and experimentation by a broader community.

The NAIRR as a Catalyst for More System Testing (Topic C)

We view testbeds as a critical and underinvested facet of AI R&D, and echo the Task Force’s findings that robust testing can spur innovation, assess the safety of AI systems, and draw on the expertise of a broader range of academic disciplines. Of the Task Force’s definition of testbeds, Anthropic has primarily focused on the use of “data sets and frameworks that support evaluation,” though we view the NAIRR as a much-needed resource to also host testing environments. In addition to the benefits outlined in Findings 4-10 – 4-13, a rigorous testing framework can assist model developers, the broader research community, and the general public in evaluating AI systems for performance and safety. Specifically, testbeds can help with:

- **Internal Benchmarking**: When building models, developers can use testbeds to understand how particular implementation decisions affect model performance and safety.
- **External Benchmarking**: Developers, the broader research community, and the general public can use testbeds to compare the relative performance and safety features of models from different organizations side-by-side.
- **Resource Forecasting**: Developers can use testbeds to more predictably anticipate future resource needs and investment decisions.
- **Understanding Societal Impacts**: Developers, the broader research community, and the general public can use testbeds to evaluate models for important societal impacts (e.g. fairness, bias, and alignment with human values).

Finding 4-11 accurately notes that testbeds can increase equitable involvement in AI research by including less well-funded institutions – this can be expanded to include the opportunity that testbeds create for bringing in new perspectives and academic disciplines. We believe the most promising AI research will come from cross-disciplinary collaboration, not only in model development, but in model testing and validation, as well. **Unlike model development, investigating existing large-scale models is a relatively low-cost way to involve the diverse and varied expertise of the academic community.** For example, the Centre for the Governance

of AI estimated that a research project examining bias in GPT-3 required less than \$100 in compute resources¹⁰. At that cost, the NAIRR could potentially support thousands of researchers across the U.S. in carrying out impactful and socially-relevant work.

Recommendation 4-19 of the Interim Report (cataloging AI testbeds in the NAIRR) will help the broader research community identify and access the methods of testing that, as of today, are spread across the research literature, development community, and open source code repositories. Supporting this effort with a dedicated, full-time staff (as proposed in Recommendation 4-17) is an excellent use of NAIRR funds and will help keep these resources up-to-date and accessible. A catalog of testbeds will provide an excellent foundation for a more expansive vision of the NAIRR in measuring and assessing AI systems. **By classifying the kinds of testbeds within the catalog, the NAIRR staff will be able to identify gaps in existing testing methods and highlight areas for additional investment.** These findings can then be used to encourage researchers to develop *new* testbeds in particular domains, which the NAIRR can fund through its research proposal review process.

Implement the NAIRR in a Staged Approach (Topic B)

We are encouraged by the ambitious goals stated by the Task Force and the potential for the NAIRR to become the leading shared infrastructure for AI R&D. To ensure its success, we recommend the NAIRR development staff triage desired capabilities and services, and implement them in a staged way. While the NAIRR will ultimately provide for a range of infrastructure, data, and support options to a varied group of stakeholders, we recommend starting with the minimum viable product (“MVP”) version of a NAIRR that can help accelerate research at the outset.

As we suggested in our previous comment to the Task Force, **the first phase of the NAIRR should prioritize resources that are critical to research and readily-available**¹¹. These include: access to cloud infrastructure, a repository of easily accessible datasets, a catalog of existing testbeds, and sufficient funding for engineering support staff. While this first iteration could support a wide variety of research efforts, **it is essential that one of the initial goals of the NAIRR be to establish infrastructure that can support the provisioning of at least one large experiment**, where “large” involves more than 100 accelerator chips (e.g. GPUs) running in parallel to train a single machine learning (ML) model¹². Because infrastructure behaves

¹⁰ Anderljung, M., Heim, L., & Shevlane, T. (2022, April 11). *Compute Funds and Pre-trained Models*. Centre for the Governance of AI. <https://www.governance.ai/post/compute-funds-and-pre-trained-models>

¹¹ Anthropic. (2021, October). *Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses*. ai.gov. <https://www.ai.gov/rfi/2021/86-FR-39081/Anthropic-NAIRR-RFI-2021.pdf>

¹² To put this in perspective, a handful of recent research projects from industry developers used several times the number of GPUs proposed here. Meta recently released a model trained on over 900 A100 GPUs (“OPT: Open Pre-trained Transformer Language Models” - <https://arxiv.org/abs/2205.01068>), while a research effort from OpenAI

differently at large scales relative to small scales, funding dedicated DevOps and systems administration staff at the outset will help get researchers up and running with large-scale projects. Running such workloads early in the life of the NAIRR can serve to quickly validate the effectiveness of the infrastructure and identify areas for further improvement.

Alongside measures of infrastructure efficacy, the NAIRR staff should continuously evaluate other performance and usability metrics to inform subsequent phases of development. Staff may consider regular user polls or automated metrics that capture whether the NAIRR expands access to AI R&D resources (e.g. geographic representation of grant recipients), whether the NAIRR meets the resource needs of researchers (e.g. percentage of requests that are met), and its overall usability (e.g. number of days to launch a project after funding). With these insights, the NAIRR staff can then work to add in complementary resources such as hybrid infrastructure solutions, previously unreleased government datasets, and more complex testing environments.

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 and encourage more equitable participation in AI R&D. Anthropic firmly supports the goals of the NAIRR and sees it as a tremendous opportunity to support critical and underfunded research into this transformative technology. Advancements made possible by the NAIRR could add to the long legacy of foundational research in academia that eventually powers technological innovation across the U.S. economy.

We urge the Task Force to put forward an ambitious vision for the scope of the NAIRR and funding, drawing inspiration from other “Big Science” infrastructure investments in fields such as physics and astronomy. We also recommend the Final Report include explicit proposals for compute allocations that can support industry-scale model development, as well as funding and development opportunities for new AI testbeds. We appreciate the opportunity to share our feedback and are eager to continue supporting the Task Force as it prepares its Final Report.

used over 700 V100 GPUs (“Video PreTraining (VPT): Learning to Act by Watching Unlabeled Online Videos”- <https://cdn.openai.com/vpt/Paper.pdf>).

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Centre for the Governance of AI (GovAI)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



Comments on the interim report of the National Artificial Intelligence Research Resource Task Force

June 30, 2022

Lennart Heim
Research Scholar
Centre for the Governance of AI

Markus Anderljung
Head of Policy
Centre for the Governance of AI

About the Centre for the Governance of AI (GovAI)

The Centre for the Governance of AI (GovAI) is a nonprofit based in Oxford, UK, with a US presence. It was founded in 2018, initially as part of the Future of Humanity Institute at the University of Oxford, before becoming an independent research organization in 2021. GovAI's mission is to build a global research community, dedicated to helping humanity navigate the transition to a world with advanced AI. More information at [governance.ai](https://www.governance.ai).

Our comments

We welcome the opportunity to submit comments on the interim report of the National AI Research Resource's task force and look forward to future opportunities to input on the NAIRR. We offer the following submission for your consideration.

Key Recommendations

We recommend that the NAIRR:

Provides researchers with access to pre-trained models by

1. providing infrastructure that enables API-based research on large pre-trained models and guards against misuse
2. allowing researchers to use their NAIRR compute budget to do research on models accessed through an API
3. exploring ways to incentivize technology companies, academic researchers, and government agencies to provide structured access to large pre-trained models through the API

Addresses the risks stemming from AI models developed with NAIRR resources by

4. implementing a tiered access approach to compute provision, where access to larger amounts of compute comes with additional review requirements

Recommendations on topic (c)

(c) NAIRR resource elements and capabilities. Including data, government datasets, compute resources, testbeds, user interface, and educational tools and services. (Chapter 4 of the report)

Also [available on the GovAI blog](#).

Compute funds and pre-trained models

One of the key trends in AI research over the last decade is its growing need for computational resources. Since 2012, the compute required to train state-of-the-art (SOTA) AI models has been doubling roughly every six months¹. Private AI labs are producing an increasing share of these high-compute SOTA AI models², leading many to worry about a growing compute divide between academia and the private sector³. Partly in response to these concerns, there have been calls for the creation of a National AI Research Resource (NAIRR)⁴. The NAIRR would help provide academic researchers with access to compute, by either operating its own compute clusters or distributing credits that can be used to buy compute from other providers⁵. It would also further support academic researchers by granting them access to data, including certain government-held datasets.

We argue that for the NAIRR to meet its goal of supporting non-commercial AI research⁶, its design must take into account what we predict will be another closely related trend in AI R&D: an increasing reliance on large pre-trained models, accessed through application programming interfaces (APIs). Large pre-trained models are AI models that require vast amounts of compute to create and that can often be adapted for a wide array of applications. The most widely applicable of these pre-trained models have recently been called foundation models⁷, because they can serve as a “foundation” for the development of many other models. Due to commercial considerations and concerns about misuse⁸, we predict that private actors will become increasingly hesitant to allow others to download

¹ [Sevilla et al., 2022](#)

² According to [Sevilla et al., 2022](#), every AI system that has set a new record for compute consumption since 2016 has been produced by a private lab.

³ [Ahmed & Wahed 2020](#); [Ganguli et al. 2022](#)

⁴ [Etchemendy & Li 2020](#)

⁵ [Ho et al. 2021](#)

⁶ [Ho et al. 2021](#)

⁷ [Bommasani et al. 2021](#)

⁸ [Brundage et al. 2018](#)

copies of these models. We instead expect these models to be accessible primarily through APIs, which allow people to use or study models that are hosted by other actors. While academic researchers need access to compute and large datasets, we argue that they will also increasingly require API access to large pre-trained models. (Lohn & Musser have made similar claims.⁹) The NAIRR could facilitate such access by setting up infrastructure for hosting and accessing large pre-trained models and inviting developers of large pre-trained models (across academia, industry, and government) to make their models available through the system. At the same time, they could allow academics to use NAIRR compute resources or credits to work with these models.

The NAIRR has an opportunity, here, to ensure that academic researchers will be able to learn from and build upon some of the world's most advanced AI models. Importantly, by introducing an API, the NAIRR could provide structured access¹⁰ to the pre-trained models so as to reduce any risks they might pose, while still ensuring easy access for research use. API access can allow outside researchers to understand and audit these models, for instance identifying security vulnerabilities or biases, without also making it easy for others to repurpose and misuse them.

Concretely, we recommend that the NAIRR:

1. provides infrastructure that enables API-based research on large pre-trained models and guards against misuse;
2. allows researchers to use their NAIRR compute budget to do research on models accessed through an API; and
3. explores ways to incentivize technology companies, academic researchers, and government agencies to provide structured access to large pre-trained models through the API.

Signs of a trend

We predict that an increasing portion of important AI research and development will make use of large pre-trained models that are accessible only through APIs. In this paradigm, pre-trained models would play a central role in the AI ecosystem. A large portion of SOTA models would be developed by fine-tuning¹¹ and otherwise adapting these models to particular tasks. Commercial considerations and misuse concerns would also frequently prevent developers from granting others access to their pre-trained models, except through APIs. Though we are still far from being in this paradigm, there are some early indications of a trend.

⁹ [Lohn & Musser 2022](#)

¹⁰ [Shevlane 2022](#)

¹¹ Fine-tuning describes the process of improving the performance of a pre-trained model on a specific task by training it on a task-related dataset.



Particularly in the domain of natural language processing, academic research is beginning to build upon pre-trained models such as T5, BERT, and GPT-3.¹² At one of the leading natural language processing conferences in 2021, EMNLP¹³, a number of papers were published that investigated¹⁴ and evaluated¹⁵ existing pre-trained models. Some of the most relevant models are accessible only or primarily through APIs. The OpenAI API for GPT-3, announced in June 2020¹⁶, has been used in dozens of research papers¹⁷, for example investigating the model's bias¹⁸, its capabilities¹⁹, and its potential to accelerate AI research by automating data annotation²⁰. Furthermore, Hugging Face's API interface has been used to investigate COVID-19 misinformation²¹ and to design a Turing test benchmark for language models²².

At the same time, in the commercial domain, applications of AI increasingly rely on pre-trained models that are accessed through APIs. Amazon Web Services, Microsoft Azure, Google Cloud²³, and other cloud providers now offer their customers access to pre-trained AI systems for visual recognition, natural language processing (NLP), speech-to-text, and more. OpenAI reported that its API for its pre-trained language model GPT-3 generated an average of 4.5 billion words per day²⁴ as of March 2021, primarily for commercial applications.

Five underlying factors in the AI field explain why we might expect a trend towards academic research that relies on large pre-trained models that are only accessible through APIs:

- Training SOTA models from scratch requires large amounts of compute, precluding access for actors with smaller budgets. For instance, PaLM²⁵ – a new SOTA NLP model from Google Research – is estimated to have cost between \$9 and \$23M to train.²⁶ The training compute cost of developing the next SOTA NLP model will likely be even greater.
- In comparison, conducting research on pre-trained models typically requires small compute budgets. For instance, we estimate that a recent paper investigating

¹² [Raffel et al. 2019](#); [Devlin et al. 2018](#); [Brown et al. 2020](#)

¹³ [EMNLP Conference 2021](#)

¹⁴ [Wolfe & Clasikan 2021](#)

¹⁵ [Elazar et al. 2021](#)

¹⁶ [OpenAI 2020](#)

¹⁷ [Google Scholar search](#)

¹⁸ [McGuffie and Newhouse 2020](#)

¹⁹ [Kohler and Daniel 2021](#)

²⁰ [Wang et al. 2021](#)

²¹ [Wahle et al. 2021](#)

²² [Uchendo et al. 2021](#)

²³ [Google 2022](#), [AWS 2022](#), [Microsoft Azure 2022](#)

²⁴ [OpenAI 2020](#)

²⁵ [Chowdhery et al. 2022](#)

²⁶ [Heim 2022](#)



anti-muslim bias in GPT-3²⁷ likely required less than \$100 of compute.²⁸ Developing new SOTA models by fine-tuning or otherwise adapting “foundation models” will also typically be dramatically cheaper than developing these models from scratch.

- The developers of large pre-trained models are likely to have strong incentives not to distribute these models to others, as this would make it both more difficult to monetize the models and more difficult to prevent misuse.
- Given the right infrastructure, it is significantly easier for researchers to use a pre-trained model that is accessed through an API than it is for them to implement the model themselves. This would enable user-friendly and secure access to the NAIRR (which is discussed in recommendation 4-20 of the interim report²⁹). Implementing large models, even for research purposes, can require significant engineering talent, expertise, and computing infrastructure. Academics and students often lack these resources.
- Academics may increasingly aim their research at understanding and scrutinizing models, as this is important scientific work and plays to academia's comparative advantage.

We discuss these factors in detail in [our blog post](#).

How the NAIRR could provide access to pre-trained models

We offer a sketch of how the NAIRR could provide access to pre-trained models in addition to data and compute, illustrated in the figure below. First, it would create a platform for hosting and accessing pre-trained models via an API. The platform should be flexible enough to allow researchers to run a wide range of experiments on a range of models. It should be capable of supporting fine-tuning, interpretability research, and easy comparison of outputs from multiple models. The API should allow researchers to interface with both models hosted by the NAIRR itself and models hosted by other developers, who may often prefer to retain greater control over their models.

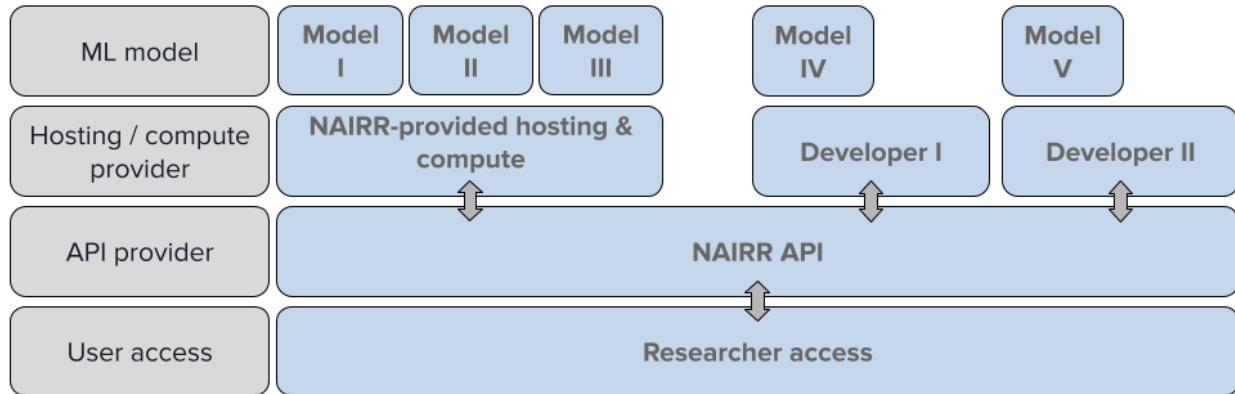
Second, researchers would be allowed to use their NAIRR compute budgets to run inferences on the models. We recommend that researchers be allowed to use their budgets for this purpose even if the model is hosted by an organization other than the NAIRR.

²⁷ [Abid et al. 2021](#)

²⁸ The authors probably used less than 10,000 prompts of around 20 tokens and received 10,000 outputs of around 20 tokens. This sums up to a total cost of around \$24 via the OpenAI Davinci API ([\\$0.06 per 1,000 tokens](#)). This would be cheaper if using a less powerful version of GPT-3 or when the inference is self-hosted.

²⁹ “*Recommendation 4-20: To help realize its vision, the NAIRR must provide secure and user-friendly access to integrated services, resources, data, and training materials.*”

In recommendation 4-13 of the interim report, “three levels” of the NAIRR compute resources are suggested.³⁰ The proposed API would be part of the third and highest level – providing access to pre-trained models for a wide range of users.



An illustration of how the NAIRR could provide API access to large pre-trained models.

The biggest challenge will likely be securing access to pre-trained models from developers across industry, academia, and government. In some cases, developers might be motivated to provide access by a desire to contribute to scientific progress, the prospect of external actors finding issues and ways to improve the model, or a belief that it might improve the organization’s reputation. The NAIRR could also create an expectation that models trained using NAIRR compute should be accessible through the platform. Access to particularly high-stakes government models in need of outside scrutiny could also potentially be mandated. Additionally, the NAIRR could consider incentivizing government agencies to provide API access to some of their more impactful models in exchange for access to compute resources or data (similar to a Stanford HAI proposal regarding data access³¹).

Encouraging private actors to make their models accessible through the platform may be especially difficult. In some cases, companies may provide model access as a means to build trust with their consumers. They may recognize that the public will be far more trusting of claims concerning the safety, fairness, or positive impacts of their AI systems if these claims are vetted by outside researchers. For example, Facebook and Twitter have recently created APIs that allow outside researchers to scrutinize company data in a privacy-preserving manner.³² Further, the NAIRR could consider offering compensation to developers for making their models available via the API. Developers may also be particularly concerned about risks to intellectual property, something that can be assuaged by the NAIRR upholding high cybersecurity standards.

³⁰ “Recommendation 4-13: Software leveraged for NAIRR compute resources should span three “levels” to support a broad user base.”

³¹ [Ho et al. 2021](#)

³² [TechCrunch 2021](#); [Twitter 2022](#)

Crucially, the API should also be designed to thwart model misuse, while still ensuring easy access for research use. Multi-purpose models trained with NAIRR resources could be used maliciously, for instance by criminals, propagators of misinformation, or autocratic governments around the world. Large language models could, for example, significantly reduce the cost of large-scale misinformation campaigns³³. The NAIRR should take measures to avoid models trained with publicly funded compute being put to such uses. Misuse could be reduced by introducing a tiered access approach, as suggested in the Stanford HAI report³⁴ for datasets hosted on the NAIRR. For instance, researchers might get easy access to most models but need to apply for access to models with high misuse potential. Further restrictions could then be placed on the queries or modifications that researchers are allowed to make to certain models. In addition, API usage should be monitored for suspicious activity (e.g. the generation of large amounts of political content).

Helping academic researchers share their models

An appropriately designed API could also solve a challenge the NAIRR will face as it provides compute and data for the training of large-scale models: academic researchers will likely want to share and build on models developed with NAIRR resources. At the same time, open-sourcing the models may come with the risk of misuse in some cases. By building an API and agreeing to host models itself, the NAIRR can address this problem: it can make it easy for researchers to share their models in a way that is responsive to misuse concerns.

Academics are significantly more likely to voluntarily make their models available via the API than private developers of SOTA models with a profit motive. As such, the NAIRR could start by focusing on providing infrastructure for academic researchers to share their models with each other, thereby building a proof-of-concept, and later introducing additional measures to secure access to models produced in industry and across government. Eventually, the NAIRR could set a standard — enabling a vibrant and growing AI ecosystem, as proposed in recommendation 4-22³⁵, while maintaining critical security needs.

Conclusion

By building API infrastructure to support access to large pre-trained models, the NAIRR could produce a number of benefits. First, it could help academics to scrutinize and understand the most capable and socially impactful AI models. Second, it could cost-effectively grant researchers and students the ability to work on frontier models. Third, it could help researchers to share and build upon each other's models while also avoiding risks of misuse. Concretely, we recommend that the NAIRR:

³³ [Weidinger et al. 2021](#); [Buchanan et al. 2021](#)

³⁴ [Ho et al. 2021](#)

³⁵ "Recommendation 4 -22: The NAIRR should embrace standards, including de facto standards, and best-of-breed open-source solutions whenever possible to ensure a vibrant, growing AI ecosystem."

1. provides infrastructure that enables API-based research on large pre-trained models and guards against misuse;
2. allows researchers to use their NAIRR compute budget to do research on models accessed through an API; and
3. explores ways to incentivize technology companies, academic researchers, and government agencies to provide structured access to large pre-trained models through the API.

Recommendations on topic (d)

(d) System security and user access controls. (Chapter 5 of the report)

We recommend that the NAIRR task force implements a tiered access scheme to computational resources (in short *compute*) — similar to the recommendation for the access to sensitive and private data.³⁶ Since compute is a finite, rivalrous resource, the NAIRR will have to make difficult decisions about how it is allocated. Such decisions should be based on many factors, including scientific merit and practicability. Importantly, it should also be based on the extent to which the researchers adhere to responsible AI practices, e.g. foreseeing and preventing potential risks the model could impose. The more compute a project is allocated, we argue, the greater care should be taken by the NAIRR and the researchers to reduce risks and spread the benefits of the system.

Chapter 5 of the interim report outlines the security and user access to the NAIRR. We welcome and support the outlined recommendations for protecting sensitive and private data. Nonetheless, the recommendations do not sufficiently address the potential risks stemming from AI systems created with resources by the NAIRR. As many scholars have argued, AI systems can pose a variety of risks and should undergo an extended review process before their creation and potential publication.³⁷ As an example, systems that can read and write can substantially impact daily life,³⁸ and their surprising and unpredictable capabilities³⁹ warrant an extensive review and monitoring process. The NAIRR should facilitate and enforce responsible development and disclosure of these powerful AI systems. The NAIRR should become leader in the co-development of these guidelines and enforce them for research conducted using the NAIRR — helping to set the standards for responsible AI.⁴⁰

³⁶ See Recommendation 4-9 (p.4.5) and 5-3 (p.5.3) of the interim report.

³⁷ Brundage et al. 2018 "[The malicious use of artificial intelligence: Forecasting, prevention, and mitigation.](#)"

³⁸ See a recent discussion by leading AI labs on "[Best Practices for Deploying Language Models](#)"

³⁹ Ganguli et al. 2022 "[Predictability and Surprise in Large Generative Models](#)"

⁴⁰ As outlined in the Executive Summary: "*The NAIRR can set the standard for responsible AI research through the design and implementation of its governance processes.*" (p.iii)

Why use compute as a proxy for the potential impact of a system?

The compute used to train an AI system is a particularly useful metric when considering what level of responsible AI practices should be demanded. Firstly, the performance of machine learning models tend to scale with compute.⁴¹ State-of-the-art models across domains, such as PaLM, AlphaFold, GPT-3, have one thing in common: they use a large amounts of compute.⁴² For example, it took more than 64 days across thousands of chips to train PaLM with an estimated cloud computing cost of \$9M to \$23M.⁴³ While the performance of a system also scales with the amount and quality of data,⁴⁴ there are no agreed-upon metrics of data quality that could be used for this purpose. Secondly, the performance of an AI system is a useful proxy of its potential impact, both positive and negative. The more capable the system, the more uses it can be put to, and the more important it is that it is developed and deployed responsibly.

Other metrics should also be considered. For example, the NAIRR could introduce stricter requirements for AI systems used in particularly high risk domains such as health care or biometric identification. However, it is often difficult to predict the downstream impacts of an AI system or research contribution.⁴⁵ An AI system developed for one use can often be put to others: an AI drug discovery tool could be repurposed to design biochemical weapons or other toxic substances.⁴⁶ Further, high-compute models trained today – and likely trained using NAIRR resources – tend to be general AI systems⁴⁷, where it is even more challenging to predict the uses or even capabilities.⁴⁸ As such, considering only, for example, the uses a model will be put to is not sufficient.

What could a compute-based tier of responsible AI practices look like in practice?

Researchers or projects receiving small amounts of compute could be subject to minimal or no responsible AI requirements. They could be required to submit a description of what they plan to use the resources for and sign an agreement saying that they will adhere to a NAIRR code of conduct. The NAIRR should also do spot checks to see if compute is being used for the intended purpose.

⁴¹ Kaplan et al. 2020 "[Scaling Laws for Neural Language Models](#)"; Hofman et al. 2022 "[Training Compute-Optimal Large Language Models](#)"

⁴² Sevilla et al. 2022 "[Compute Trends Across Three Eras of Machine Learning](#)"

⁴³ Chowdhery et al. 2022 "[PaLM: Scaling Language Modeling with Pathways](#)"; Heim 2022 "[Estimating PaLM's training cost](#)"

⁴⁴ Model size (number of parameters) and number of data samples are linear correlated with the amount of compute. However, it's independent of the quality of data.

⁴⁵ Prunkl et al. 2021 "[Institutionalizing AI Ethics via Broader Impact Statements](#)"

⁴⁶ Urbina et al. 2022 "[Dual use of artificial-intelligence-powered drug discovery](#)"

⁴⁷ Bommasani et al. 2021 "[On the Opportunities and Risks of Foundation Models](#)"

⁴⁸ Ganguli et al. 2022 "[Predictability and Surprise in Large Generative Models](#)"

At the higher end, a number of requirements could be imposed on the project. As a starting point, such developers of such models could be required to adhere to the forthcoming NIST AI Risk Management Framework⁴⁹, in addition to an extended review process and policies around the future publication and usage.

Importantly, a number of measures could be taken to ensure that potential risks from the system are identified and mitigated. Identifying such risks can be hard, as it is difficult to predict what tasks a general model will perform well at,⁵⁰ and because the eventual impacts of the system depends on how it gets incorporated into larger sociotechnical systems. This could be done by requiring external audits or red team exercises. It could also be done by giving initial access to a few dozen researchers and giving them a bias/safety bounty if a flaw in the model is identified.⁵¹

Potential risks from the system could be addressed in the development phase, ensuring that the models are sufficiently accurate, fair, robust, aligned, interpretable and the like. Some risks can be addressed via appropriate deployment strategies or "structured access".⁵² For example, particularly general and capable models could be made available to a wide audience using an API, with monitoring and filters to prevent misuse.

Deciding on the thresholds for the responsible AI tiers will be a challenging task. A useful starting reference might be a fraction of the compute used for the final training run of state-of-the-art AI systems, likely measured in FLOPs. The final training run compute is commonly reported by researchers and scholars have tracked it over time.⁵³ While those reports only address the compute used for the final training run, not the complete development process, this can be used as a lower bound for the required compute for developing AI systems. Further, the compute thresholds should likely differ depending on the type of system or application domain, as e.g. SOTA models in protein folding require much less compute than those in natural language processing.

⁴⁹ NIST 2022 [AI Risk Management Framework](#)

⁵⁰ Ganguli et al. 2022 "[Predictability and Surprise in Large Generative Models](#)"

⁵¹ For a description of these tools and more, see Brundage et al., 2021. "[Toward Trustworthy AI Development](#)"

⁵² Shevlane 2022 "[Structured access: A paradigm for safe AI deployment](#)"

⁵³ Sevilla et al. 2022 "[Compute Trends Across Three Eras of Machine Learning](#)"; Amodei & Hernandez 2018 "[AI and Compute](#)"

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Consumer Reports

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



June 30, 2022

National AI Research Resource Task Force
Attn: Ms. Wendy Wigen, NCO, NITRD Program
2415 Eisenhower Avenue
Alexandria, Virginia 22314

Re: Request for Information Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force

Dear Members of the National AI Research Resource Task Force:

Consumer Reports (CR) writes today in response to the Request for Information Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force. 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. 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. This is important not just for AI advancement, but also provides researchers with the tools to identify and call-out harm that can be caused by AI. 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

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.

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 into everyday products and daily life, it is important that its development be democratized and accessible to all in order to mitigate harmful effects.

Vision for the NAIRR

While providing computing resources to researchers and small businesses is a worthy goal, we primarily advocate for NAIRR being an ethical resource that sets industry standards for best practices. We recommend that NAIRR should work with other federal agencies that enforce civil rights law to release guidelines on how industry should go about designing, testing, and deploying algorithms in different sectors to mitigate harm. Currently, there is a lack of industry standards in terms of quality of training data, privacy protections to mitigate identification of individuals that could be identified by data, accuracy rates of algorithms prior to deployment, testing and maintenance of algorithms, and requirements for independent auditing of algorithms (particularly ones with sensitive applications). These are all areas that AI researchers, companies, and the public could benefit from more guidance. AI has the potential to roll back much of the civil rights protections that have been afforded to us, and providing computational and data resources to researchers and companies on its own will not solve this problem. We will elaborate what we hope to see NAIRR contribute to AI standards and ethics in the following sections.

NAIRR resource elements and capabilities

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 perform research on and release guidelines regarding the potential misuse or misapplication of AI, preferably in conjunction with other federal agencies that enforce civil rights law. We applaud the interim report particularly for "*Recommendation 4-8: The NAIRR should establish a value ecosystem around data that can be used for AI.*" For example, the use of pseudoscience and physiognomy are on the rise in AI applications; some companies claim that their AI can do things that are not necessarily possible or substantiated by science.² NAIRR should make clear why certain uses of AI are harmful 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

² Narayanan, Arvind. "How to Recognize AI Snake Oil."
<https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf>

privacy-protecting methods that allow for careful examination of how civil rights can be potentially impacted by AI without disclosing anyone's personal information.

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 independent 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.

Privacy, civil rights, and civil liberties requirements

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.³ 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.

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,

³ 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>.

public-interest researchers attempting to audit or reverse engineer potentially harmful algorithms are not able to do so without higher-quality training data.

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.

However, combining datasets from different sources and ensuring that datasets are comprehensive across different demographics could lead to an erosion of privacy since data can be used to point to specific individuals. It is important that NAIRR takes this into account when providing data to the public or researchers; processes should be put in place to de-identify and anonymize data to the extent possible and only provide researchers with data necessary to complete their projects.

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; and, we hope this initiative can be used to provide much needed guidance for ethical standards in industry. Thank you for your consideration.

Sincerely,
Nandita Sampath
Policy Analyst

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Data Foundation

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

June 30, 2022

RE: Feedback on the findings and recommendations put forward in the NAIRR Task Force's Interim Report

To Whom It May Concern:

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 is America's premier voice on data policy, advocating for responsible policies to make government data high-quality, accessible, and usable. Ensuring reasonable, responsible, and ethical practices are implemented in legislative and administrative activities is a priority for the Data Coalition, thus we work to promote strategies for meaningful artificial intelligence (AI) advancements in government.

The Data Coalition applauds the NAIRR Task Force for the publication of its interim report. The report is a productive step in the long journey toward making access to data for AI research equitable, secure, and effective for use across sectors. We fully support the goals, objectives, and overall vision of the NAIRR. The following response expands on three areas identified in the request for feedback, emphasizing aspects from the interim report that we consider critical to include in the final report, and highlights areas where additional elaboration may be helpful.

B. Establishment and sustainment of the NAIRR. Including agency roles, resource ownership and administration, governance and oversight, resource allocation and sustainment, and performance indicators and metrics. (Chapter 3)

The identified management and governance structure of the NAIRR as a federated cybersecurity ecosystem is conducive to ensuring the proper independence, oversight, and transparency of a national AI R&D tool. Incorporating a tiered access system and performance measures will both encourage use and highlight the value of a shared data infrastructure.

The more NAIRR data is used, the more there will be opportunities to identify flaws and solutions, strengthening the functions and quality of the data and research products stemming from a NAIRR. Recommendations such as 3-13 to incentivize use, and recommendation 3-17 to evaluate the impact of NAIRR use, are important aspects to ensure NAIRR data is functioning as envisioned and contributing meaningfully to future AI R&D.

Further, as seen in recommendation 3-16, calling for external evaluation is important for transparency, accountability, and ensuring effectiveness of a NAIRR. Without identifying external evaluators, evaluations may be limited to a small number of evaluators close to the

project, potentially influencing the findings or direction of evaluations that should be objectively conducted.

Also relevant to sustainability and governance, while the report acknowledges the need for Congress to fund NAIRR through appropriations and suggests the management entity can explore other revenue streams for long-term sustainability, it would be beneficial to provide a cost estimate. Including an estimate of how much funding NAIRR will need from Congress can facilitate further advocacy around NAIRR implementation to ensure it has the funding to achieve the Task Force's vision.

C. NAIRR resource elements and capabilities. Including data, government datasets, compute resources, testbeds, user interface, and educational tools and services. (Chapter 4)

The report highlights that “[i]ncreased access and diversity of perspectives would, in turn, lead to new ideas that would not otherwise materialize and set the conditions for developing AI systems that are inclusive by design.” We see this as a priority for the NAIRR and we support the explicit recommendations expanding access to and use of data resources across the government.

Seen in recommendations 4-24 through 4-26, the Task Force's provisions for software, training, and educational resources to support a diverse set of users with varying levels of proficiency is important to achieving the vision of a NAIRR. Supporting a career pathway – rather than a career pipeline – can help expand and better ensure inclusion, access, and equity in a data workforce. Incorporating these educational resources and activities are key factors for a NAIRR to broaden the contributions to AI R&D.

The search and discovery portal for data is also an important function of the NAIRR, seen in recommendation 4-8 as well as 2-3. Providing a single access portal that incorporates existing data repositories from across the government, rather than duplicating efforts, can expand data access and use, leverage existing data capacity, improve overall government efficiency, and facilitate contributions to AI R&D that may have previously been excluded. We would like to reiterate the need for a streamlined, standardized, searchable, and accessible portal to achieve the full potential of a NAIRR as identified in the goals and objectives. We also encourage the task force to connect with other single application portal initiatives taking place within the federal statistical system, as well as the Advisory Committee on Data for Evidence Building.

Similarly, coordination of open data plans, among other efforts, seen in Recommendation 4-11 is crucial to leverage current advancements in AI R&D and government-wide data use and data sharing. A NAIRR that aligns with the Foundations for Evidence-based Policymaking Act requirements will bolster the federal data infrastructure and can serve as an important model that may be applied in other research areas. All of these steps to ensure accessible and open data will improve transparency and help promote trustworthy AI initiatives.

E. Privacy, civil rights, and civil liberties requirements. (Chapter 6)

The role of the NAIRR management entity to implement steps to ensure transparency, access for diverse users, and proper oversight as well as demonstrate how research using the NAIRR is being reviewed, approved, and performed are all critical to protecting civil rights and liberties. We support these steps and believe by expanding upon NAIRR data curation processes and identifying standards for documenting bias both within the research process and the data itself, the NAIRR can further bolster responsible use of AI.

Overall, we agree with with the interim report that a “NAIRR would transform the U.S. national AI research ecosystem by strengthening and democratizing foundational, use-inspired, and translational AI R&D in the United States.” As the Task Force prepares its final report, we encourage continued emphasis on expanding data access and use, developing a search and discovery portal, and ensuring data and research processes are without bias or inaccuracies.

Thank you for the opportunity to provide comments on this very important issue, and we hope to continue to support your efforts to support ethical and useful AI advancements in government. Please contact me at corinna@datafoundation.org if you have any questions or would like to discuss the Data Coalition’s interest in this matter further.

Sincerely,

Corinna Turbes
Managing Director
Data Coalition

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Greg Dreifus, Luis Videgaray Caso

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

Democracy and Competition in AI: **A proposal for a decadal survey in advanced computing**

Greg Dreifus* and Luis Videgaray Caso**

* Massachusetts Institute of Technology¹

Introduction

“There are some who maintain that democracy cannot cope with a new technique of government developed in recent years by a few countries that deny the freedoms that we maintain are essential to our democratic way of life. That I reject.” - Franklin Roosevelt, 1939

While the United States is in many ways a global leader in AI technology, its ability to harness AI for the public good is hindered by how it invests in this crucial area. Spending on computing innovation is driven mostly by large corporations. This means that the vision for AI’s future, what problems will be solved and what values technology reflects, will be driven by the values of private sector actors and the incentives to which they are accountable. This is occurring amidst a backdrop of a ballooning global competition between democracies and authoritarians to shape the 21st century, the core institutions that govern it, and the values, priorities, and norms of societies around the world. Central to this competition is the race for AI supremacy. Who develops AI systems will harness a powerful tool in steering the course of everyone’s geopolitical future.

The United States needs a plan to harness AI’s research and development in an inclusive, disciplined manner. Such a plan should be focused towards a concrete vision, with goals set forth about what the United States wants AI to look like as it evolves over time. The United States has in fact already shown that successful long-term, consistent spending in large-scale research and development (R&D) is possible in democracy, and therefore such planning is not necessarily a competitive advantage of autocracy. America has achieved success in implementing sustained planning in the space sciences using a system known as the decadal surveys. The institutional framework of the decadal process collates the vision of the space science community to plan and ultimately execute lengthy and expensive space science missions, even as political priorities shift over time as a result of the democratic process. While there are notable differences between the space sciences and AI, we suggest adapting this decadal framework into a Computing and Artificial Intelligence Decadal Plan (CAIDP) to ensure the time consistency and long-term vision of government spending in the field of advanced computing and AI.

The Decadal Precedent: How to spend more public money on computing research and development?

¹ We submit this as individuals. Our views are not representative of the Massachusetts Institute of Technology, the MIT Sloan School of Management, the MIT Schwarzman College of Computing, or the AI Policy Forum.

The United States can learn from well established precedents in developing a disciplined long-term plan for public AI investments aligned with democratic values. Namely, it can learn from the decadal model. The United States has shown that the decadal model has worked successfully at NASA. Although NASA has experienced certain challenges over its history with cost overruns, the end results of its R&D policy are ultimately a paragon of scientific achievement, which includes landing a new reconnaissance rover on Mars as recently as 2021, and its leadership continues despite recent tumults in U.S. politics. This goal driven scientific research and development took root in the United States during World War II. The Manhattan Project and similar war efforts led to an enormous amount of innovation at the time. Immense government resources funded the development of the atomic bomb, radar, and penicillin. At the end of the war, White House science advisor and former MIT Vice President Vannevar Bush published his seminal “Science: The Endless Frontier” that laid out the principles of prolonged public R&D during peacetime.²

The framework established by Bush evolved into newer dimensions at the dawn of the space race. The National Academies of Sciences, Engineering, and Medicine founded the Space Science Board (SSB) in 1958, one year after the launch of Sputnik.³ In 1965, the SSB convened a group of experts in Woods Hole, Massachusetts to develop a plan for the coming decade to map the trajectory of the space sciences.^{4,5} The Woods Hole conference resulted in the first in a sequence of so-called decadal reports that survey the foremost specialists in the space community to understand the most promising directions in their fields and make recommendations for government investment.

The decadal survey is a robust, intensive, and formalized strategic plan organized for the space sciences community. After the initial decadal report was released in 1965, numerous follow-on reports were published throughout the 1970s and 1980s, including “Strategy for Exploration of the Inner Planets (1977-1987),”⁶ “Strategy for the Exploration of Primitive Solar-System Bodies--Asteroids, Comets, and Meteoroids (1980-1990),”⁷ and “A Strategy for Exploration of the Outer Planets (1986-1996).”⁸ The current surveys focus on five subject areas: Astronomy and Astrophysics, Solar and Space Physics, Biological and Physical Sciences, Earth Science and Applications from Space, and Planetary Sciences.^{9,10} Decadal surveys provide long term agendas on expensive research missions and act as guideposts to the various governmental

² <https://www.nsf.gov/od/lpa/nsf50/vbush1945.htm>

³ <https://www.nationalacademies.org/ssb/about>

⁴ <https://www.hq.nasa.gov/office/pao/History/SP-4214/ch3-6.html>

⁵ <https://www.nap.edu/catalog/12410/space-research-directions-for-the-future>

⁶ <https://www.nap.edu/catalog/12379/strategy-for-exploration-of-the-inner-planets-1977-1987>

⁷

<https://www.nap.edu/catalog/12372/strategy-for-the-exploration-of-primitive-solar-system-bodies-asteroids-comets-and-meteoroids>

⁸ <https://www.nap.edu/catalog/12345/a-strategy-for-exploration-of-the-outer-planets-1986-1996>

⁹ <https://science.nasa.gov/about-us/science-strategy/decadal-surveys>

¹⁰

<https://www.nap.edu/initiative/committee-for-the-decadal-survey-on-biological-and-physical-sciences-in-space>

agencies, including NASA, NSF, and the Geological Survey, as well as Congressional committees and offices that issue appropriations.¹¹ The decadal process is still operational, and the most recent decadal is “Vision and Voyages for Planetary Science in the Decade 2013-2022.”¹² The decadal process tradition is still working even under the current difficult political environment.

The planning for ongoing decadal reports entails extensive reviews that allow for the deliberation of as many stakeholders as possible in the outcome of the process. The survey’s formulation lasts for two years. The outcome explains where a field of study should dedicate specific resources, what facilities should be built, and how much large-scale projects are likely to cost. Congress then uses the book-long decadal reports as guidance to write line items in the government’s budget to invest in the recommendations of the decadal committees. From this process, plans are made to build satellites, telescopes, and other foundational research enterprises, and to conduct celestial exploration to foreign bodies like Mars. The procedures are carried through on a bipartisan basis and overlap between consecutive administrations and Congressional turnover. The longevity of decadal planning makes them robust against the vicissitudes of election cycles and even against increasing American political polarization.

The relative success of the decadal surveys makes them a compelling model for other complex and long-term research endeavors, including artificial intelligence and advanced computing. A computing decadal survey could specify the most critical and challenging hurdles in AI development and specify the resources needed to overcome them in a way to best improve the welfare of the nation. A report could help Congress target, with money, those offices that need it to achieve these goals, and it would mitigate the likelihood of inefficiencies in project management. Perhaps most important, an approach inspired by the current decadal surveys would streamline a durable vision where the research should head, circumventing short sighted, politically motivated, or ineffectual benchmarks towards achieving innovation in the computing field.

Other efforts have been made to develop U.S. AI strategies to varying degrees over the years. The U.S. Global Change Research Program seeks to “coordinate federal research and investments in understanding the forces shaping the global environment, both human and natural, and their impacts on society,” but it does not focus on AI.¹³ The High Performance Computing Initiative strategizes around a narrow subset of AI systems but is not centered around a value-based mission. The National Security Commission on Artificial Intelligence laid out an extensive AI strategy but is much broader than the remit suggested here, encompassing issues like human talent, national security strategy, technology transfer and more. Moreover, it does not lay out how to update its own suggestions through an institutional process. We suggest

¹¹ <https://www.nap.edu/catalog/18434/lessons-learned-in-decadal-planning-in-space-science-summary-of>
¹²

<https://solarsystem.nasa.gov/resources/598/vision-and-voyages-for-planetary-science-in-the-decade-2013-2022/>

¹³ <https://www.globalchange.gov/about>

a sustained, institutional framework oriented around values and a mission would be novel and constructive.

There are some noteworthy differences between the space sciences and AI that the Office of Science and Technology Policy and the National Science Foundation should consider in putting together a Computing and Artificial Intelligence Decadal Plan (CAIDP). Artificial intelligence is a general purpose technology, which can be implemented in multifaceted fields, sectors, and applications. The space sciences and space missions largely require hardware-based systems that come with known price tags, narrow purposes and design features, and that are already being discussed by the relevant scientific community. For this reason, decadal survey budgets are largely known *a priori* (even if inaccurately). The National Academy of Sciences can therefore conduct the decadal survey by tailoring the agenda of a narrow subset of the space science community to a very focused agenda. The breadth of AI distinguishes how an AI decadal-like process would be executed.

Despite these differences, the space science decadal can still serve as inspiration for a CAIDP in many notable ways. Precisely because AI is diffuse, broad, and general, it would benefit from a focused plan with concrete vision and mission. Rather than spending money on a hodgepodge of projects across fields, the government would be able to orient the most critical and desired features for the future of AI and dedicate resources to align AI with democratic values. To efficaciously implement a decadal-long Computing and Artificial Intelligence Decadal Plan (CAIDP), the following principles should be observed.

Guiding Principles

A CAIDP is an opportunity to build an inclusive long-term vision about the development and deployment of AI in America that includes everybody, a vision that represents the interests not only of technology companies but also of workers and other stakeholders that would be strongly affected by AI. It is important that a CAIDP aligns with the interests of diverse segments of the population across demographics, economics, and political power with a concrete plan considering and engaging with the needs of marginalized and historically underrepresented groups and minorities. Therefore, a CAIDP should commit to guiding principles that can allow it to fully realize its intended mission.

1. **Interagency Engagement:** The charter of such a decadal should help coordinate the numerous federal departments that depend on the future of computing technology. These include agencies explicitly dedicated to scientific innovation like the National Science Foundation, the National Institute of Health, the Department of Energy's Office of Science, NASA, NOAA, U.S. Geological Survey, and NIST. They also include major government actors in research at the Department of Defense.
2. **Defining the Scope:** A computing decadal survey should specify the most critical and challenging hurdles in AI development and specify the resources needed to overcome them in a way to best improve the welfare of the nation.

3. Stakeholder Engagement: Relevant stakeholders cut across a varied segment of society, and a successful decadal survey will capture the perspective and consensus of all informed and impacted parties.
 - a. Private Sector: Currently, a substantial portion of computing research is carried out by the private sector, which would require any computing decadal commission to engage industry in a way not historically as necessary in the space sciences community. In the formation of the decadal committees, separate planning committees could be formed to represent major technology companies, small and medium sized companies, and startups. Each subcommittee could be tasked to pick technical representatives to participate on their behalf in the decadal formation process, creation of the statement of task, and the decadal process itself.
 - b. Unions: AI should be developed and deployed in a way that is not only about automation but also about creating new jobs and sharing the prosperity that AI can create. Worker representation should therefore influence the vision for AI systems and innovation.
 - c. Non-governmental Organizations: Computing research requires a degree of ethical and societal consideration, pursuant to the weighty consequences of AI innovation, that is perhaps different than developments in many other technologies. As such, various non-technical stakeholders, like NGOs, should be included in the drafting of the decadal survey.
 - d. Academia and the Research Community: Thought leaders, computer scientists, social scientists, and philosophers in the researcher community, in academia, and at public laboratory facilities should be central to orchestrating the CAIDP.
 - e. Policymakers: The CAIDP process should also include and engage with elected officials at all levels of government (federal, state, and local) to both ensure the CAIDP reflects the democratic instincts of the citizenry and to engage with the needs of federalized American governance.
4. Guiding Ethics and transparency: Extra guardrails and procedural steps could go a long way in preventing undue influence by factional interests, political biases, corrupting lobbying practices, or financial incentives. Transparency is key to ensuring all participants are driven by scientific goals and not self-serving agendas. The experts who ultimately conduct and implement the decadal should stay alert to participants' conflicts of interests or pursuit of private agendas.
5. Strategic Innovation Mapping: It is helpful to break down the possible dimensions of governmental research and development efforts. Along one axis, the government can invest in fundamental research, applied research, or translational research. Along another axis, the government can invest in research that would be suboptimally done by the private sector left alone, most often in basic research, or governments can support commercial research in direct partnership with the private sector, driven by their needs, resources, and networks. For each of these dimensions of investment, the government can take a goal oriented approach to research and development, as it has done in the original development of the atomic bomb and for space exploration, or not, allowing open ended and exploratory R&D to expand the societal body of knowledge.

6. Planning Inception Process: The initiation of the CAIDP process will be pivotal to the survey's success, and attention must be paid to establishing a quality foundation for the process to proceed.
7. Adapting Over Time: The value in the decadal is that it is a consensus building tool wielded by experts who listen to the opinions of relevant stakeholders. By design, it will evolve to the demands of current events in the computing field while focusing on long term, big picture problems.

Proposal for a Computing and Artificial Intelligence Decadal Plan

1. Establish an Initiation Committee in the National Academy of Sciences (NAS) to define the scope and list of stakeholders to engage in drafting the CAIDP. This process should define what vision is preferred for the future of AI, what mission it wants to achieve, and what diverse segments of society (including which sectors) are central to achieving these ends.
2. Because of NAIIA's passage, the CAIDP can commence by springboarding off the NAIRR. A NAIRR can provide recommendations for the CAIDP, including its scope and structure, by making recommendations to the NAS for the Initiation Committee and subsequent CAIDP structure and execution.
3. Establish a Strategic Commission on AI Planning (SCAIP) at the NAS, which will run the sustained CAIDP process given the resources and staff at its disposal. The SCAIP should be composed of subcommittees dedicated identifying the critical topic areas in computing sciences and roadmaps for investments in each.
4. In selecting what a CAIDP might encompass, the SCAIP can look across the diversity of computing sciences and other fields that directly intersect with AI, taking into consideration where government investments can have the greatest public benefit and where private industry underinvests. It will be critical in the formation of the CAIDP to identify the meta questions of how the decadal should orient itself to yield the optimal public outcome.
5. Congress should pass a law codifying the structure of the CAIDP and SCAIP at the NAS, building off the work of the NAIIA and NAIRR.

Conclusion

Technological advances in AI have already upended many aspects of society, and AI is already having a destabilizing effect in our world. The ascendancy of authoritarianism has also exploited AI systems to suppress freedom, diminish democracy, and empower dictatorship. There is no guarantee that AI will ultimately have a net positive upshot. Perhaps most critically, without a tangible vision and without leadership, there is no guarantee computing innovation will yield a beneficial outcome just because it happens within a democratic country. The future of AI systems is a choice to be made, determined by what research initiatives are picked, who runs them, how money is spent, and what plans are implemented. The Computing and Artificial Intelligence Decadal Plan, building upon the architecture put in place by a NAIRR, would compel the US democracy to decide, within the parameters of democratic norms and structures,

the future that AI creates for us all. By implementing a CAIDP, the United States, as a leader of global democracies, can undertake an open, transparent, inclusive, and democratic process to deliver consistent investment to a sector crucial for the future of the world. In so doing, the United States can prove Franklin Roosevelt right by leading the vanguard of democratic government towards the ends of outcompeting dictatorship, this time through a long term and inclusive vision of AI that embraces the principles of democracy and shared prosperity..

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Electronic Privacy Information Center (EPIC)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

COMMENTS OF THE ELECTRONIC PRIVACY INFORMATION CENTER

to the

White House Office of Science and Technology Policy and National Science Foundation

Regarding the

Implementation of the Interim Report of the National AI Research Resource

87 FR 31914

June 30, 2022

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.¹ EPIC has a long history of promoting transparency and accountability for information technology.²

EPIC recommends the White House Office of Science and Technology Policy (“OSTP”) and the National Science Foundation (“NSF”)³ center human rights, civil rights, and thoughtful procurement and reduce reliance on and federal funding of private sector AI systems in their implementation of the National AI Research Resource Task Force Interim Report (“NAIRR Report”)⁴ In October 2021, EPIC urged the NAIRR 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.

¹ EPIC, *About EPIC* (2022), <https://epic.org/epic/about.html>.

² EPIC, *Algorithmic Transparency* (2022), <https://www.epic.org/algorithmic-transparency/>; EPIC, *Algorithms in the Criminal Justice System* (2022), <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*, Federal Trade Commission (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 Educational, Scientific and Cultural Organization (“UNESCO”) (Mar. 15, 2018), 5-6, [https://epic.org/internetuniversality/EPIC_UNESCO_Internet_Universality_Comment%20\(3\).pdf](https://epic.org/internetuniversality/EPIC_UNESCO_Internet_Universality_Comment%20(3).pdf).

³ Hereinafter to referred to as “agencies”

⁴ National Science Foundation, Science and Technology Policy Office, *Request for Information on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force*, Federal Register, 87 FR 31914 (June 30, 2022) <https://www.federalregister.gov/documents/2022/05/25/2022-11223/request-for-information-rfi-on-implementing-initial-findings-and-recommendations-of-the-national>

EPIC renews these points to the agencies as they work to implement NAIRR’s interim report. EPIC recommends that the agencies set defined minimum standards for algorithmic transparency and accountability and set strict limits on the use of NAIRR to develop certain harmful applications like emotion recognition.

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 careful adoption of algorithms.⁵ EPIC has advocated for transparency and accountability internationally, litigating cases against the U.S. Department of Justice to compel production of documents regarding “evidence-based risk assessment tools”⁶ and against the U.S. Department of Homeland Security to produce documents about a program to assess the probability that an individual commits a crime.⁷ In 2018, EPIC and leading scientific societies petitioned the U.S. Office of Science and Technology Policy to solicit public input on U.S. Artificial Intelligence Policy.⁸ EPIC submitted comments urging the National Science Foundation to adopt the UGAI and to promote and enforce the UGAI across funding, research, and deployment of U.S. AI systems.⁹ EPIC has also 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 adequate regulation to protect individuals.¹⁰

⁵ See e.g. EPIC v. DOJ (D.C. Cir.) (18-5307), <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*, Department of Housing and Urban Development (Oct. 18, 2019), <https://epic.org/apa/comments/EPIC-HUD-Oct2019.pdf>; Testimony of EPIC, Massachusetts Joint Committee 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 Committee on Commerce, Science & Transportation (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 Science and Technology Policy (Apr. 4, 2014), <https://epic.org/privacy/big-data/EPIC-OSTP-Big-Data.pdf>.

⁶ EPIC, *EPIC v. DOJ (Criminal Justice Algorithms)*, <https://epic.org/foia/doj/criminal-justice-algorithms/>.

⁷ See *Id.* and EPIC, *EPIC v. DHS (FAST Program)* <https://epic.org/foia/dhs/fast/>.

⁸ EPIC, Petition to OSTP for Request for Information on Artificial Intelligence Policy (July 4, 2018) <https://epic.org/privacy/ai/OSTP-AI-Petition.pdf>.

⁹ EPIC, Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan, National Science Foundation, 83 FR 48655 (Oct. 26, 2018) <https://epic.org/apa/comments/EPIC-Comments-NSF-AI-Strategic-Plan-2018.pdf>.

¹⁰ Comments of EPIC, *Solicitation of Written Comments by the National Security Commission on Artificial Intelligence*, 85 Fed. Reg. 32,055, National Security Commission on Artificial Intelligence (Sep. 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. 1825, 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 Commission 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 Commission (Sep. 10, 2020) <https://epic.org/apa/comments/EPIC-EU-Commission-AI-Sep2020.pdf>.

To establish necessary consumer safeguards, EPIC has filed FTC complaints against HireVue,¹¹ an employment screening company, and Airbnb,¹² the rental service that claims to assess risk in potential renters based on an opaque algorithm. EPIC has also filed a petition with the FTC for a rulemaking for AI in Commerce.¹³

Agencies Implementing NAIRR’s Interim Plan Must Strengthen Oversight to Ensure That Protections for Individual Rights Aren’t Overtaken by Commercial Interests

EPIC urges the agencies to focus on building capacity for oversight and strictly purposeful procurement and away from aiding development for purpose of international competition. Throughout the report, the focus is on increased capacity to create new AI systems¹⁴, but without parallel focus to control harmful AI and with insufficient recognition about the lack of accuracy or control. Several of the recommendations from the NAIRR report could help strengthen protections for people subject to AI, but the devil is in the details, and EPIC provides specific ways in which the recommendations with protective potential must be implemented to be meaningful.

The agencies must prioritize building oversight capacity for an infrastructure to perform independent and thorough audits and impact assessments for both commercial and governmental AI, as well as effective ways to communicate the findings of those oversight mechanisms, and triggers that take unacceptably risky AI off the market.

The strongly worded set of suggestions found in Recommendation 3-4¹⁵, that the operations of the NAIRR should live outside of the government to not be stymied by burdensome procurement regulations is troubling. This necessarily sets up incentives away from the protection of fundamental rights and toward further regulatory capture. Again, the government should build capacity for oversight and careful development, not defer it, and the interest should be strictly in the hands of a party without independent profit motives.

Recommendation 3-6¹⁶ and 3-7¹⁷ are illustrative of the report’s insufficient focus on meaningful protection. EPIC believes that increased flexibility in contracting with private

¹¹ Complaint and Request for Investigation, Injunction, and Other Relief, *In re HireVue* (Nov. 6, 2019), https://epic.org/privacy/ftc/hirevue/EPIC_FTC_HireVue_Complaint.pdf.

¹² Complaint and Request for Investigation, Injunction, and Other Relief, *In re Airbnb* (Feb. 27, 2019), https://epic.org/privacy/ftc/airbnb/EPIC_FTC_Airbnb_Complaint_Feb2020.pdf.

¹³ *In re: Petition for Rulemaking Concerning Use of Artificial Intelligence in Commerce*, EPIC (Feb. 3, 2020) <https://epic.org/privacy/ftc/ai/EPIC-FTC-AI-Petition.pdf>.

¹⁴ NAIRR Interim Report, *supra* note 10 2-2.

¹⁵ NAIRR Recommendation 3-4: The day-to-day operations of the NAIRR should be managed by an independent, non-governmental entity with dedicated, expert staff.

¹⁶ NAIRR Recommendation 3-6: The NAIRR management entity should have flexibility in contracting, partnering, or entering other agreements with the private sector, with appropriate government oversight.

¹⁷ NAIRR Recommendation 3-7: The NAIRR management entity should be explicitly charged with addressing diversity, equity, inclusion, and accessibility (DEIA) issues related to NAIRR-supported AI R&D.

companies beyond what the government already does would be harmful to the quality of AI that comes out of NAIRR. EPIC supports the centering and concretizing of DEIA principles in all AI R&D, but a vague charge of addressing is not enough. There need to be meaningful controls as well as prohibitions on what types of AI is developed, and the oversight and accountability capacity to limit AI harm of those communities.¹⁸

Recommendations 3-7 and 3-8¹⁹ must include specific prohibitions and limitations on what tools can be developed. Without this, Recommendation 3-9²⁰ calling for governing charters and policies amounts to administrative theater.

The agencies should be extremely cautious in their implementations of NAIRR's Recommendations that incentivize data collection, including sensitive data collection, and include strong data minimization requirements to all projects.

EPIC commends and supports a broad implementation of Recommendation 3-10.²¹ Access to NAIRR resources should be contingent on review, clear use policies must include prohibitions, and shared information must be made transparent to those subjected. Recommendations 3-7 – 3-9 must be read into the use policies of 3-10.

¹⁸ Algorithmic harm is rampant and felt hardest by marginalized communities. See e.g., Larry Hardesty, *Study finds gender and skin-type bias in commercial artificial-intelligence systems*, MIT News (Feb. 11, 2018), <http://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212>; Melissa Hamilton, *The Biased Algorithm: Evidence of Disparate Impact on Hispanics*, 56 Am. Crim L. Rev. 1553 (2019), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3251763; Megan T. Stevenson & Christopher Slobogin, *Algorithmic Risk Assessments and the Double-Edged Sword of Youth*, 96 Wash. U.L. Rev. 681 (2018), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3225350; See, e.g., Muhammad Ali, Piotr Sapiezynski, Miranda OgenOgen, Aleksandra Korolova, Alan Mislove, Aaron Rieke, *Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes*, arXiv:1904.02095 (Apr 3, 2019); GAO summary of Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan. "Dissecting racial bias in an algorithm used to manage the health of populations," *Science*, vol. 366, no. 6464 (2019), pp. 447-453 | GAO-21-519SP; Tom Simonite, *How an Algorithm Blocked Kidney Transplants to Black Patients*, (Oct. 26, 2020) <https://www.wired.com/story/how-algorithm-blocked-kidney-transplants-black-patients/>; Amit Datta, Anupam Datta, Michael Carl Tschantz, *Automated Experiments and Privacy Settings A Tale of Opacity, Choice, and Discrimination*, arXiv:1408.6491v2 [cs.CR] 17 (Mar. 18 2015), available at <https://arxiv.org/pdf/1408.6491.pdf>; Sheridan Wall and Hilke Schellmann *LinkedIn's job-matching AI was biased. The company's solution? More AI*, MIT Technology Review (Jun 23, 2021) <https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence/>; Daan Koklman, "F**k the algorithm?" *What the world can learn from the UK's A-level grading fiasco*, LSE Impact Blog (Aug. 2020).

¹⁹ NAIRR Recommendation 3-8: NAIRR management and administration should be governed by a formal charter and associated policies, with an executive leadership team managing day-to-day operations.

²⁰ NAIRR Recommendation 3-9: The governance policies and performance of the NAIRR should be overseen by a board of governors and complemented with mechanisms for external advice, oversight, and evaluation

²¹ NAIRR Recommendation 3-10: Access to NAIRR resources should be contingent on research project proposal review, be governed by clear use policies and user agreements, and be in compliance with relevant requirements for open sharing of research outputs.

For Recommendation 3-15²², the agencies must require definition of metrics and indicators of success grounded in established best practices to be made public, and established as part of a open and public process not controlled by industry preference. Metrics must be created to evaluate how the AI that comes out of NAIRR supported projects are protective of rights, transparent, and in line with a defined set of principles of AI. EPIC recommends the use of UGAI which has been endorsed by more than 250 experts and 60 organizations in 40 countries, as well as the Organisation for Cooperation and Economic Development AI Principles to guide R&D.²³ The UGAI comprise twelve principles:

1. Right to Transparency.
2. Right to Human Determination.
3. Identification Obligation.
4. Fairness Obligation.
5. Assessment and Accountability Obligation.
6. Accuracy, Reliability, and Validity Obligations.
7. Data Quality Obligation.
8. Public Safety Obligation.
9. Cybersecurity Obligation.
10. Prohibition on Secret Profiling.
11. Prohibition on Unitary Scoring.
12. Termination Obligation.²⁴

The OECD AI Principles²⁵ were adopted in 2019 and endorsed by 42 countries—including the United States and the G20 nations.²⁶ The OECD AI Principles establish international standards for AI use:

1. Inclusive growth, sustainable development and well-being.
2. Human-centered values and fairness.
3. Transparency and explainability.
4. Robustness, security and safety.
5. Accountability.²⁷

EPIC strongly agrees with Recommendation 3-19,²⁸ and urges the agencies to include tracking evaluations in line with the above principles, as well as accuracy rates and bias reporting

²² NAIRR Recommendation 3-15: NAIRR evaluation methods, including definition of metrics and indicators of success for the NAIRR, should be grounded in established best practices.

²³ *Universal Guidelines for Artificial Intelligence*, The Public Voice (Oct. 23, 2018), <https://thepublicvoice.org/ai-universal-guidelines/>.

²⁴ *Id.*

²⁵ *Recommendation of the Council on Artificial Intelligence*, OECD (May 21, 2019) [hereinafter *OECD AI Principles*], <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>;

²⁶ *U.S. Joins with OECD in Adopting Global AI Principles*, NTIA (May 22, 2019), <https://www.ntia.doc.gov/blog/2019/us-joins-oecd-adopting-global-ai-principles>.

²⁷ *OECD AI Principles*, *supra* note 15.

²⁸ NAIRR Recommendation 3-19: The NAIRR management entity should establish a publicly accessible platform that tracks the usage and outputs of NAIRR-supported research and the results of external evaluations.

in the publicly accessible outputs. These requirements will help legitimate each system and increase trust.

Conclusion

EPIC urge the agencies to implement the findings of the NAIRR Interim Report with strict prohibitions, reporting requirements, transparency, and keeping control for government actors. The NAIRR should not primarily be an accelerator for AI development for AI development's sake. EPIC will be happy to provide further comment in later stages of implementation.

Respectfully Submitted,

s/ Ben Winters
Ben Winters
EPIC Counsel

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Engine

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



June 30, 2022

National AI Research Resource Task Force
Attn: Ms. Jeri Hessman, NCO, NITRD Program
2415 Eisenhower Avenue
Alexandria, VA 22314
(202) 459-9683

VIA ELECTRONIC SUBMISSION

Re: Request for Information (RFI) on Implementing Initial Findings and Recommendations of the National Artificial Intelligence 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. Emerging technologies like artificial intelligence are being developed and utilized by startups to solve problems across a range of industries—from cybersecurity to agriculture and beyond. Accordingly, we appreciate the Task Force’s iterative process and opportunity to comment on the development of NAIRR and the Interim Report.

I. The NAIRR user base must be broadened to include a more diverse group of startups (Topic A).

Expanding access and lowering barriers to AI research—including for private entities like startups—are goals of the NAIRR and its original sponsors in Congress,¹ but the user base as currently envisioned may exclude many startups that should be beneficiaries and have access to the NAIRR. The Interim Report outlines that only startups “that have been awarded Federal grants via the Small Business Innovation Research (SBIR) or Small Business Technology Transfer (STTR), or other similar Federal programs for small businesses...” will have full access to the NAIRR.²

¹ The strategic objective of the NAIRR reflected in the interim report is to “strengthen and democratize the U.S. AI innovation ecosystem,” of which startups must necessarily be a part. *Envisioning a National Artificial Intelligence Research Resource (NAIRR): Preliminary Findings and Recommendations*, NAIRR Task Force (May 2022), <https://www.ai.gov/wp-content/uploads/2022/05/NAIRR-TF-Interim-Report-2022.pdf>; see also e.g., 166 Cong. Rec. H3501 (Jul. 20, 2020) (statement of Rep. Eshoo) and 166 Cong. Rec. H6932 (Dec. 8, 2020) (statement of Rep. Eshoo).

² See Interim Report, supra note 1.

The SBIR and STTR programs, while helpful for startups that receive them, can be insular, struggle with representation, and have long timelines that might not best suit startups that could benefit from access to the NAIRR. SBIR recipients broadly have raised this concern in the past. For example, Neil Ray, an SBIR awardee and the Founder & CEO of San Ramon, California-based Raydiant Oximetry notes that while integral for his company, the program’s “timeline doesn’t work for startups that need to act quickly,” and discusses the insular, academic nature of the grant program that could exclude some founders without access to those networks.³ And the application-based method of (less than full) access to the NAIRR for those who are not federal grantees would likely face similar challenges for the timeline necessary in order to be useful for fast-moving startups.

Federal grant programs, including SBIR and STTR, have historically struggled to be demographically representative, meaning their use as grant of access to the NAIRR may exclude those for whom the resource is designed to democratize access to AI research. While reflective of larger societal issues with diversity in business ownership (and noting quality of demographic data may be lessened due to self-reporting), just 13.3% of all SBIR grants from 2011-2018 were awarded to companies led by women.⁴ Similarly low numbers of grants go to minority-owned companies.⁵

And companies that have said the NAIRR would benefit them would likely not have access based upon their non-grantee status. For example, the startup Beehero, who responded to the NAIRR RFI in Fall 2021 and relayed such a message, does not appear to be a recipient of SBIR, STTR or other federal grant based on public information.⁶

Access to the NAIRR is important for startups, because, as Engine highlighted in our initial comments to the Task Force, AI research and development can be prohibitively expensive and out of reach for many startups operating on resource-limited budgets.⁷ This makes the establishment of the NAIRR an 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. But the NAIRR will only be successful in meeting those ends if it includes a sufficiently broad group of users—including those in the startup ecosystem—and access is granted in a way that is not cumbersome and sufficiently quick for startups.

³ #StartupsEverywhere Profile: Neil P. Ray, MD, Founder & CEO, Raydiant Oximetry, Engine (May 27, 2022), <https://www.engine.is/news/startupseverywhere-sanramon-ca-raydiantoximetry>.

⁴ Jenny Servo, et al., *Women’s Inclusion in Small Business Innovation Research & Small Business Technology Transfer Programs*, Americas Seed Fund, 16 (Aug 2020), https://cdn.www.nwbc.gov/wp-content/uploads/2020/08/11124006/Women-In-SBIR-Report_NWBC_Final_2020-08-07.pdf.

⁵ See, e.g., *Assessment of the SBIR and STTR Programs at the National Institutes of Health* National Academies of Sciences, Engineering, and Medicine, 118, <https://nap.nationalacademies.org/read/26376/chapter/6#118>.

⁶ *Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource*, BeeHero, (Sept 30, 2021), <https://www.ai.gov/rfi/2021/86-FR-39081/BeeHero-NAIRR-RFI-2021.pdf>.

⁷ Engine, RFI Response: National AI Research Resource (NAIRR), (Sept. 1, 2021), <https://www.ai.gov/rfi/2021/86-FR-39081/Engine-NAIRR-RFI-2021.pdf>.

II. Options for broader inclusivity of startup participants in the NAIRR (Topic A).

To broaden inclusivity of the NAIRR and enable startup participation, the task force may consider leveraging entities with local knowledge of startups and, separately, develop a vetting procedure to allow full access to the NAIRR that is timely while recognizing legitimate security concerns.

Many startup communities are anchored by entrepreneur support organizations (ESOs), like co-working spaces, incubators, or accelerators—most of which are organized as not-for-profit. These organizations are trusted by startups and disseminate important information and resources to them. At the same time, ESOs also have deep knowledge of the companies they work with and support. Vetting ESOs to clear access based upon their local knowledge of startups would reduce the number of organizations in need of scrutiny by administrators of the NAIRR by distributing responsibilities for allowing access, and would speed access for startups.

University innovation centers serve as anchoring ESOs for some startup ecosystems. While they tend to be more insular than other ESOs, given NAIRR's nexus to universities, it may be worthwhile to consider leveraging these centers to facilitate startup access.

Finally, the NAIRR Task Force could conceive of its own unique vetting process for granting full access to the resource. While recognizing legitimate security concerns, the process must be quick, accessible, and understandable, so as not to discourage startups from seeking the resources at the outset.

* * *

Thank you for the opportunity to provide these comments regarding access to the NAIRR. Engine remains committed to engaging with the Task Force and stands ready to be a resource in thinking through ways to support innovators and startups across the country advancing AI research.

Engine
700 Pennsylvania Ave. SE
Washington, DC 20003

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Hugging Face

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



Hugging Face
29 June 2022

20 Jay St
Suite 620
New York, NY 11201

Hugging Face Comments on Implementing Findings from the National Artificial Intelligence Research Resource Task Force

Hugging Face commends the National Artificial Intelligence Research Resource (NAIRR) Task Force on its interim report and we offer recommendations to further shape innovation for good, responsible artificial intelligence (AI). The following comments are informed by our experiences as an open platform for state-of-the-art (SotA) AI systems, working to make AI accessible and broadly available to researchers for responsible development. Comments are organized by Interim Report Chapter, with more granular recommendations italicized below. If a section or chapter is not highlighted, we do not have specific, actionable feedback.

About Hugging Face

Hugging Face is a community-oriented company based in the U.S. and France working to democratize good machine learning. We are an open-source and open-science platform hosting machine learning models and datasets within an infrastructure that supports easily processing and analyzing them; conducting novel AI research; and providing educational resources, courses, and tooling to lower the barrier for all backgrounds to contribute to AI.

a. Vision for the NAIRR (Chapter 2)

Appoint Technical and Ethical Experts as Advisors

Technical experts with a track record of ethical innovation should be prioritized as advisors. In order for NAIRR to drive innovation in a responsible direction, it must craft a diverse external advisory body with interdisciplinary expertise. As part of Recommendation 2-3, technical and ethical experts can calibrate NAIRR on not only what is technically feasible, implementable, and necessary for SotA systems, but also on how to avoid exacerbating harmful biases and other malicious uses of AI systems. [Dr. Margaret Mitchell](#), one of the most prominent technical experts and ethics practitioners in the AI field and Hugging Face's Chief Ethics Scientist, is a natural example of an external advisor.

c. NAIRR resource elements and capabilities (Chapter 4)

Resource (Model and Data) Documentation Standards

NAIRR-provided standards and templates for system and dataset documentation will ease accessibility and function as a checklist. This standardization should ensure readability across audiences and backgrounds; documentation should be robust for researcher and developer information, clearly have examined and reported ethical considerations, and be easily consumable for a nontechnical audience. [Model Cards](#) are a vastly adopted structure for documentation that can be a strong template for AI models. [Datasheets](#) are the strong parallel, also widely adopted but for datasets.

Making ML Accessible to Interdisciplinary, Non-Technical Experts

In addition to being user-friendly, NAIRR should work toward encompassing the many critical expertises in AI that may not have the advanced technical knowledge to leverage the provided resources. Combining recommendations 4-20, 4-23, 4-24, and 4-26, NAIRR should provide education resources as well as easily understandable interfaces and low- or no-code tools for all relevant experts to conduct complex tasks, such as training an AI model. For example, Hugging Face's [AutoTrain](#) empowers anyone regardless of technical skill to train, evaluate, and deploy a natural language processing (NLP) model. Hugging Face [Tasks](#) and [Spaces](#) enable anyone to build and engage with tasks, from prompting the image-generative model [Dall-E Mini](#) to a [Wikipedia Assistant](#) that can answer open-ended questions based on Wikipedia content.

d. System security and user access controls (Chapter 5)

Guardrails for Open Source Science: Monitor for High Misuse and Malicious Use Potential

In addition to cybersecurity, a key aspect of preventing harm from dual-use AI systems and resources is to evaluate and monitor for high potential for harm. Datasets and models created with the intent to harm or that overrepresent harmful content should be closely monitored or gated to prevent bad actor access. Harm must be defined by NAIRR and advisors and continually updated, but should encompass egregious and harmful biases, political disinformation, and hate speech. NAIRR should also invest in legal expertise to craft [Responsible AI Licenses](#) to take action should an actor misuse resources.

e. Privacy, civil rights, and civil liberties requirements (Chapter 6)

Empowering Diverse Researcher Perspectives via Accessible Tooling and Resources

Tooling and resources must be available and accessible to different disciplines as well as the many languages and perspectives needed to drive responsible innovation. Implementing Finding 6-1 (Engaging diverse stakeholders) is critical for incorporating the expertise and viewpoints of the many stakeholders affected by AI systems. This means at minimum providing resources in multiple languages, which can be based on the most spoken languages in the U.S. The [BigScience Research Workshop](#), a community of over 1000 researchers from different disciplines hosted by Hugging Face and the French government, is a good example of empowering perspectives from over 60 countries to build one of the most powerful open-source multilingual language models. Our platform encourages many groups to curate datasets and evaluations in their native or fluent languages makes for stronger, representative science.

Conclusion

NAIRR aligns closely with Hugging Face's mission to democratize AI in a responsible direction. We look forward to supporting this initiative as the Task Force works toward its final report and are eager to contribute to implementation.

Respectfully,

Irene Solaiman

Hugging Face

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

IBM

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

June 30, 2022

Attn: Jeri Hessman
National Coordination Office
Networking and Information Technology Research and Development
2415 Eisenhower Avenue
Alexandria, VA 22314

RE: IBM RFI Response: National AI Research Resource Interim Report [87 FR 31914]

Dear Ms. Hessman,

IBM appreciates the opportunity to comment on the Request for Information (RFI) on [Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force](#). IBM strongly supports development of the NAIRR, and many of the initial recommendations made in the interim report.

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. IBM co-created the [COVID-19 High Performance Computing Consortium](#), where 43 members have carried out more than 100 projects. 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.

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. A federated, hybrid cloud enabled computing resource – an accessible and easy-to-use hybrid- and multi-cloud computing platform built on open architecture that amalgamates various public clouds, private clouds, and on-premises resources to create a single, flexible compute infrastructure.
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.

Designed correctly, the NAIRR will be a pervasive, easily accessible federation of advanced computing resources, combined with a shared data infrastructure and tooling, that would bolster American leadership in AI research. IBM appreciates the opportunity to comment on the interim report and looks forward to future engagements.

Sincerely,

Dr. Dario Gil
Senior Vice President and Director of IBM Research

**IBM Response to the Request for Information on [Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force](#)
White House Office of Science and Technology Policy (OSTP)
National Science Foundation (NSF)**

B. Establishment and sustainment of the NAIRR. Including agency roles, resource ownership and administration, governance and oversight, resource allocation and sustainment, and performance indicators and metrics.

Agency Funding, Roles, and Responsibilities (Recommendations 3-1 to 3-3)

IBM supports recommendations 3-1 to 3-3 of the interim report. Adopting a federated approach for implementation, deployment, and administration of the NAIRR will allow agencies to maximize their collective investments and build collaboration frameworks between agencies and people. Furthermore, a federated approach would promote the sharing of diverse data sets amongst agencies and endow the NAIRR with the freedom and flexibility to grow over time.

Agencies should make new and existing federally funded (or federally owned) cyberinfrastructure, including data sets, computational resources, software and services, and testbeds, available via the NAIRR, expanding the scope and scale of the NAIRR as resources and funding permit. For example, a federated approach allows the NAIRR to incorporate 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 by enforcing open hybrid cloud architectures, and application programming interfaces.

Crucially, a federated approach would allow the NAIRR management entity to channel the perspectives of various federal stakeholders into a cohesive strategy. Therefore, IBM concurs with the report's recommendation that Congress should fund the NAIRR through appropriations to a collection of federal agencies representing AI stakeholders.

Ownership and Administration (Recommendations 3-4 to 3-7)

IBM supports recommendations 3-4 to 3-7 of the report, but suggests focusing on a specific management approach outlined in recommendation 3-4, which states the day-to-day operations of the NAIRR "could be an FFRDC, a university, a contractor, a non-profit organization, an institute, a consortium, or another such entity." Given the complex and long-term nature of the task assigned to the NAIRR, IBM recommends that it be incubated as a Federally Funded Research and Development Center (FFRDC). An FFRDC could 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. Importantly, the FFRDC model would allow multiple stakeholders (and government agencies) to rally around the NAIRR's shared goal of boosting access to AI R&D. And the FFRDC model would ensure interoperability for communities of discovery that work collaboratively to solve common challenges, such as climate change and the research challenges that will be tackled by the newly created [Advanced Research Projects Agency for Health](#) (ARPA-H).

IBM concurs with the report's recommendation that management of the NAIRR will require "a permanent and diverse staff focused on resource provisioning, managing core operations and system components – including

cataloguing external resources providing user support, and overseeing security operations.” Since the NAIRR will need to be continually updated to meet the growing and diverse needs of the AI R&D community, a dedicated staff will also be well positioned to incorporate new cyberinfrastructure resources and capabilities provisioned by the private sector and delivered through the NAIRR user access portal. To support a broad and diverse AI community (and uptake by as many researchers as possible), NAIRR management should ensure that principles of diversity, equity, inclusion, and accessibility are incorporated as it builds and operates the resource.

NAIRR Performance Indicators and Metrics (Recommendations 3-15 to 3-20)

IBM supports recommendations 3-15 to 3-20 of the interim report. To guide the development and sustainment of the NAIRR over an extended time horizon, IBM supports recommendation 3-15 on adopting a “logic model” to track progress throughout the lifecycle of the NAIRR, including its technical specification, development, testing, rollout, use, maintenance, and eventual upgrade. To foster transparency and trust in the resource, IBM supports recommendation 3-16 on the use of independent, external evaluators to assess the performance of the NAIRR, and communicate their results to Congress, sponsoring federal agencies, and the public.

IBM supports the instantiation of four levels of performance indicators outlined in recommendation 3-17, including resource investments, measures of resource usage/activities, measures of outputs, and the impact of the NAIRR and the research it enables. IBM supports the NAIRR management entity being equipped with sufficient funds to carry out data collection and evaluation (recommendation 3-18), and with this information being published in a standardized and publicly available manner to track usage and outputs from NAIRR supported research (recommendation 3-19). In its [October 2021](#) RFI response, IBM recommended that the NAIRR should adhere to [FAIR guiding principles](#) for scientific data management and stewardship. Finally, IBM supports efforts by the NAIRR management entity to use evaluation to adjust and update the goals, functions, and capabilities of the NAIRR (recommendation 3-20).

C. NAIRR resource elements and capabilities. Including data, government datasets, compute resources, testbeds, user interface, and educational tools and services.

Data (Recommendations 4-1 to 4-9)

IBM generally supports recommendations 4-1 to 4-9, and suggests areas for the task force to strengthen its approaches to data in its final report. IBM agrees that the NAIRR should coordinate a network of trusted data and compute providers to create a robust, transparent, and responsible data ecosystem. IBM supports recommendation 4-2 on following the “Five Safes” framework for safe use, which extends to the NAIRR’s security framework.

In line with recommendation 4-5 on high-quality data, 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 compliance 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.

Data acquired through the resource should be compatible 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 multiple research experiments.

To bolster recommendation 4-5, 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. Requiring interoperability, replicability and portability, in turn, ensures democratization of the resource and prevents overreliance on or lock-in by a single provider.

IBM supports the curation of a dedicated professional staff for the NAIRR, some of whom will be devoted to supporting data users, contributors, and curators. In line with recommendation 4-7, 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. The NAIRR will also need to provide tools and frameworks to enable moving and replicating data to enable computing over geographically distributed data sets.

IBM understands that potentially confidential or sensitive data will be integrated into the NAIRR, and therefore supports the creation of a tiered user access model for the resource. To the extent possible, enforcement of the security and governance requirements should be automated using available technologies.

Compute Resources (Recommendations 4-12 to 4-16)

IBM supports recommendations 4-12 to 4-16. The NAIRR can best fulfill its mission by launching a federated computing resource built on an open architecture and grounded in principles of interoperability. The open, interoperable architecture would enable the seamless integration of multivendor public and private cloud computing resources and services to create a single, unified, flexible compute infrastructure. Such an architecture would be capable of scaling and delivering compute resources in a cost-effective manner, and it would give users the flexibility to select and combine the optimal resources and computing services for their applications and to move workloads freely as circumstances change.

Such an approach will require making transitions between hybrid- and multi-cloud environments more seamless, for example, by employing a control plane that provides standardized abstractions and automation to deliver interoperability and optimal use of hybrid- and multi-cloud resources. Such interoperability would unlock new possibilities for researchers, including computational work that requires integrating data movement and computing across diverse locations and providers. In addition, commercial compute providers also provide sophisticated software stacks and user interfaces that the AI research community has widely adopted.

Crucially, recommendation 4-15 (which states that computing resources should be deployed using a phased approach) would prevent vendor lock-in and data-gravity, a phenomenon in which applications, computing, and users gravitate to a single provider over time.

Testbeds (Recommendations 4-17 to 4-19)

IBM supports recommendations 4-17 to 4-19. Testbeds can accelerate AI research by providing virtual or physical environments to test, simulate, explore, and develop AI. IBM believes that not all research conducted on the NAIRR will require running novel experiments, and the establishment of accessible testbeds that integrate computing resources, data, experiments, and evaluations for the specific AI application areas will be crucial to scaling AI R&D.

User Interface (Recommendations 4-20 to 4-23)

IBM supports recommendations 4-20 through 4-23, which state that the NAIRR should develop a central, user-friendly portal to ensure widespread user uptake and adoption. To support the modular and agile development of user interfaces, IBM supports the integration of open-source solutions and standards, where applicable.

Educational Tools and Services (Recommendations 4-24 to 4-26)

IBM supports recommendations 4-24 to 4-26. 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. Therefore, IBM supports task force recommendations aimed at building tools for community support and exchange to ensure widespread adoption of the NAIRR.

To boost adoption, IBM recommends NAIRR management focus on building peer-to-peer knowledge sharing across different scientific communities. 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.

D. System security and user access controls. (Recommendations 5-1 to 5-6)

IBM supports recommendations 5-1 to 5-6. To fulfill its quest to reach a broad and diverse set of users and boost AI R&D, the NAIRR must be secure and resilient. 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.

To this end, IBM supports the task force’s recommendation that NAIRR management create an advisory committee to recommend security standards for resources provisioned through the NAIRR. As outlined in our October 2021 [RFI response](#), such security standards may ultimately be “FedRamp-inspired.” At a minimum, the NAIRR 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. Additionally, the NAIRR must ensure compliance with additional regulations as necessary, such as HIPAA, FERPA, PCI DSS.

IBM supports recommendation 5-2, which recommends that the NAIRR adopt a zero trust architecture to ensure strong identity and access controls, including multi-factor authentication and phishing defense.

Since the NAIRR will likely be built using an open hybrid- and multi-cloud architecture, IBM supports recommendation 5-3, which calls for NAIRR management to create a tiered access model that accommodates heterogeneous security requirements. Data and compute providers should be required to attest to their security compliance and integrity prior to admission to the resource, and workloads must be restricted to data and compute resources that are allowed by their risk classifications. For example, certain 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 should 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.
6. The platform should support the integration of policy-based data governance.
7. Mixed access controls: a robust mechanism with a mix of access control models will allow for data sharing while maintaining security and privacy.

IBM strongly supports recommendations 5-4 to 5-6, which are aimed at building the NAIRR’s human capital to support the NAIRR’s system security, including providing hands on training to NAIRR staff and users; supporting the monitoring and updating of security controls; and investing in technical security experts who can ensure that the NAIRR remains in compliance with evolving security requirements.

For its part, IBM [is investing](#) in building cybersecurity resilience and the cybersecurity workforce of the future. In May 2022, IBM launched cybersecurity training programs at six Historically Black Colleges and Universities (HBCUs), which will provide cybersecurity curricula, cloud access, and an immersive learning experience to expand their capacity to develop cybersecurity talent.

IBM also supports investment in automation of compliance and security processes, including the promotion of relevant frameworks and standards, to achieve continuous compliance with cybersecurity regulations while reducing the labor cost of security operations and compliance management. IBM is actively participating in relevant standardization activities at NIST, such as OSCAL, and in the growth of open communities and projects, such as Open-SSF.

E. Privacy, civil rights, and civil liberties requirements (Recommendations 6-1 to 6-4)

IBM supports recommendations 6-1 through 6-4. Building trust in the NAIRR is paramount, and 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. By enshrining privacy, civil rights, and civil liberties requirements in the NAIRR, the resource has an opportunity to serve as an exemplar for how transparent and ethical AI R&D can be performed and scaled.

IBM strongly supports recommendation 6-2 of the report, which recommends establishing a dedicated ethics process to review all resources brought on to the NAIRR – and the research breakthroughs derived from the resource. IBM has established its own [AI Ethics board](#), which acts as a central, cross-disciplinary body to support a culture of ethical, responsible, and trustworthy AI throughout IBM. The board supports a centralized governance, review, and decision-making process for IBM ethics policies, practices, communications, research, products and services.

More broadly, IBM recommends that the NAIRR adopt [principles for trustworthy AI](#), including fairness, explainability and transparency, as they are developed by NIST as part of its [AI Risk Management Framework](#). Equipping NAIRR’s users with training on rights, responsibilities, and best practices related to privacy, civil rights, and civil liberties will boost trust in the resource, and ensure that it reaches a broad and diverse user base.

To support [bias mitigation](#), the 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. IBM also supports NAIRR's efforts to ensure transparency regarding AI data sets, common practices, and decisions that inform development of use cases, which will be instrumental in the design of new AI R&D resources and tools.

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

IEEE - USA

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

30 June 2022

To: Jeri Hessman
National Coordination Office for Networking and Information Technology Research and Development
2415 Eisenhower Avenue
Alexandria, VA 22314
NAIRR-responses@nitrd.gov

Re: RFI on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource (87 FR 31914)

IEEE-USA is pleased to submit comments in response to the findings and recommendations of the interim report by the National AI Research Resource Task Force.

The input provided below represents the expertise of the IEEE volunteer members who are living and working in the US, and who are actively conducting research and development into artificial intelligence (AI), software engineering, cybersecurity, and advanced computing. As a community of researchers, and developers, IEEE-USA strongly supports the efforts of the White House Office of Science and Technology Policy (OSTP) to ensure diversity in the R&D community, and to create a roadmap that will enable expanded access to resources, data, testbeds and associated tools for all researchers, students, and developers of AI systems.

We applaud the efforts of the Task Force and thank them for their hard work in fulfilling your 2020 Congressional mandate to create this roadmap as part of the National AI Strategy. IEEE-USA has provided specific thoughts and recommendations below. However, our overall impression is that the report could benefit from streamlining and better integration of its ideas. The report captures quite a few excellent findings, but enhanced continuity would help stakeholders better understand the impact of the report's recommendations. IEEE-USA suggests employing planning practices to the strategic objectives to streamline the roadmap and capture metrics of success. We also found that the document seemed to be duplicative in areas; this is the case for some of the recommendations regarding creation and maintenance of the resource infrastructure in Chapter 4. We also believe the document could benefit from the key recommendations and associated actions being set out as a roadmap in chart form.

The concept of explainable AI (XAI) is missing from the document. In various publications, the US government addresses the challenge of ensuring that stakeholders - both users and those impacted by AI systems - understand and trust algorithmic outcomes; see for example.^{1, 2, 3, 4} IEEE-USA recommends that NAIRR address XAI as an important research element to ensure that research outcomes produce solutions to problems that are understood by humans. Ensuring that XAI is a goal helps to replace 'black box' solutions in machine learning, where even the designers of the AI application cannot explain why the AI application reached a specific decision.

IEEE-USA is strongly supportive of the Task Force's work and believes this roadmap represents needed guidance for our AI innovation ecosystem. We thank OSTP and NSF for considering these comments and welcome further discussions with the agency on these matters. If you have questions, please do not hesitate to contact Erica Wissolik at (202) 530-8347 or e.wissolik@ieee.org.

1 DARPA: Explainable Artificial Intelligence (XAI) <https://www.darpa.mil/program/explainable-artificial-intelligence>. Accessed 29 June 2022.

2 Congressional Research Service: Artificial Intelligence: Background, Selected Issues, and Policy Considerations <https://crsreports.congress.gov/product/pdf/R/R46795>. Accessed 29 June 2022.

3 From Ethics to Operations: Current Federal AI Policy <https://atarc.org/wp-content/uploads/2022/01/Current-Federal-AI-Policy-Assessment-FINAL.pdf>, see page 8. Accessed 29 June 2022.

4 NIST: AI Fundamental Research - Explainability; <https://www.nist.gov/artificial-intelligence/ai-fundamental-research-explainability>. Accessed 30 June 2022.

IEEE-USA’s comments on the interim report:

a. A Vision for the NAIRR. (Chapter 2)

NAIRR User Base

IEEE-USA thanks the NAIRRTF for specifically mentioning students at community colleges among the users who should receive NAIRR support. Community colleges are a necessary part of the innovation economy and specifically including these students encourages diversity in the future American workforce.

Additionally, furthering the strategic objective to democratize the US AI innovation ecosystem, we believe that the NAIRR should emphasize its support for AI research at Minority Serving Institutions that serve historically underrepresented populations. For example, The University of Texas at San Antonio - which offers an excellent multidisciplinary studies degree in Artificial Intelligence - is a Minority Serving Institution due to its high number of Hispanic students.⁵ Supporting its programs, and those of similar institutions, is among the most efficient ways of pulling underserved populations into the AI workforce.

b. Establishment and sustainment of the NAIRR. (Chapter 3)

Ownership and Administration

IEEE-USA recommends that the final report include a section that clarifies the NAIRR’s management and governance structure. While the interim report alludes to the creation of a management structure and offers several alternatives, we ask that the NAIRR recommend the establishment of a shared AI resources national center management team. The structure could include a combination of the proposed alternatives to ensure that all stakeholders, in both the public and private sectors, can participate. This acknowledges the finding that no single current agency should form a new division for the reasons stated, while giving clear guidance to Congress and the Administration.

Resource Allocation and Sustainment

Recommendation 3-10

The requirement that access to NAIRR resources be contingent on research project proposal review increases the barrier of entry for students attempting to enter the field. Instead, NAIRR should look at ways to reach students in community colleges and perhaps even training programs like RAMTECH - a robotics and advanced manufacturing training center in Ohio. These stakeholders likely lack the know-how and resources to write grants.

To ensure that the research community attracts a diversity of talent, access should not be limited to only those practicing an advanced state of the art. IEEE-USA recommends that access to NAIRR resources be accessible via a tiered criteria based on level of understanding. Tiered access, where larger projects requiring more resources would have higher proposal submission requirements while smaller projects would be subjected to less stringent proposal requirements, would invite new researchers into the field, opening opportunities to less experienced professionals.

Additionally, a free “on-demand” cloud instance which serves as a sandbox (a safe environment where students can engage and learn with no financial grant) for those who may have an interest in AI systems. This could be similar to Google’s Colaboratory.

⁵ IEEE-USA Position Statement, *Artificial Intelligence: Jobs, Education, Workforce, and Diversity*, November 2021.

Ultimately, tiered access would make the NAIRR available to high level research scientists, university and community college students, and those who are just beginning to be interested in AI technologies. Since innovation can come from anywhere, the NAIRR must be open to participation from anywhere.

NAIRR Performance Indicators and Metrics

IEEE-USA recommends establishing specific measures in the roadmap that clearly explain how it will be implemented, including timelines and metrics for success. We agree with the interim report recommendation that determining value and impact of successful research programs is facilitated by clearly defined and measurable goals at the outset.

The simplest starting point for a rollout could be divided into four phases:

- Phase one - assembling the cyberinfrastructure (i.e., hardware, servers, user interface).
- Phase two - roll out platform at the community college level for debugging / beta testing. We suggest this because incomplete products at higher levels could result in loss of interest in the platform.
- Phase three - university level adoption of a platform.
- Phase four - industry level adoption.

The interim report does not address educational outcomes associated with Figure 1, Chapter 1. To ensure that the NAIRR reaches the target audience, characteristics such as gender, age, ethnicity, and program of study (i.e., engineering, physics, mathematics, biology) should be reported to demonstrate diversity in the population. These metrics could be valuable in identifying and targeting areas where the program performs weakly. This is possibly discussed in Recommendation 3-18, but the wording is unclear.

As the interim report states in the strategic objective, the NAIRR should strengthen and democratize the American AI innovation ecosystem. To successfully achieve this objective, the NAIRR must clearly measure success. We recommend using specific measures such as:

- resources sharing utilization rate,
- government branch efficiencies and operation efficiencies,
- number and type of users,
- time spent on data resources or testbeds, and
- measures of impact, both social (e.g., developed and then actually used) and academic (e.g., number of citations and influence, additional research, and other foundational work)

While the document does address strategic objectives, there is no clear strategic planning to achieve these objectives. In general, SMART (Specific, Measurable, Achievable, Relevant, and Time-bound) goals are commonly used constraints for strategic planning. These goals would clearly explain and define the strategic objective, associated measures of success, probability of success, importance of the goal being sought (the document does a great job of this), and the estimated time for completion. A source of confusion was the lack of semantic consistency when discussing the strategic objectives. Recommendations should be clearly linked to each strategic objective. SMART goals could be distributed across several sections, but all of them should be present and addressed thoroughly.

c. Resource elements and capabilities. (Chapter 4)

IEEE-USA suggests establishing a national AI resources clearinghouse as a central technical network. The National Science Foundation Network (NSFnet) could serve as a model, one where anyone who is interested can access data as well as provide information about what is available. NSFnet, which linked major universities' resources and made them available across the country, was initially restricted to government and academia. Ultimately, NSF opened access to anyone interested in and working with research or learning about computer networking in the US. The result was rapid growth of internet service providers. Establishing a clearinghouse could allow users to learn of others undertaking specific research and training utilizing NAIRR resources, with the intent of facilitating connections among disparate users in the NAIRR community.

Data

Recommendation 4-1

Recognizing that it would be impossible for the NAIRR to curate a central repository of resources, IEEE-USA suggests the design and implementation of specific metrics for quality that could be referenced by users when contributing to the central repository.

Key to implementing this recommendation is the need for a governance structure to vet candidate providers of both data and compute resources, coordinate contributions and their utilization as well as evaluate over time (e.g., annually) with suitable metrics regarding the use and scale of data sets and the computing plus network connectivity performance of contributors and any associated NAIRR resources.

Recommendation 4-5

This will be an implementation challenge. Perhaps a tiered approach may work. For example, Tier 1 - quality data suppliers exceeding X gigabytes/month or terabytes/month may access compute resources at 80 percent discount from list price services. Quality data implies that it is 'clean'.⁶ Tier 2 - quality data suppliers that supply clean data streams below the Tier 1 threshold may access compute resources at a 40 percent discount. Tier 3 - users that do not supply data pay list prices for services, with an option to reduce cost based on the type of user (e.g., K-12, college, university, not-for-profit research organizations, and collaborative research consortia involving educational institutions).

Recommendation 4-7

It is unclear where the training programs will be located. We recommend centralization and implementation through a 'school' or 'training center' operated by NAIRR staff, who would be the focal point for ensuring a qualified user community. These training centers should issue certificates of qualification for user organizations as well as individuals, on a scale based on the sophistication of the applications the user community seeks. IEEE-USA recommends that the report clarify that training programs be made available both in-person and virtually to help ensure widespread accessibility.

Recommendation 4-8

IEEE-USA recommends defining "value ecosystem" as used in this instance.

Recommendation 4-9

⁶ https://en.wikipedia.org/wiki/Data_cleansing

Without a common framework, if data providers are permitted to set the security category of data, there could be many different thresholds. This may become a management challenge. We recommend that NAIRR set the standards that data providers must use to classify their data sets. This would help eliminate the potential for varying degrees of data security access for similar data from different providers.

Compute Resources

Recommendation 4-12

This recommendation requires considerable effort to organize and execute with multiple anticipated contractual relationships between NAIRR and compute resource suppliers in government, industry, and the private sectors. When implementing, we recommend looking at similar efforts to create a federated mix of resources such as the Digital Research Alliance of Canada, an organizational vehicle that facilitates access to a variety of independently operating HPC resources.⁷ The NSFnet was an early model of sharing discreetly operated HPC resources via a shared WAN network.

NAIRR might also consider establishing relationships with exoscale compute resources for advanced AI research in partnership with National Laboratories such as Los Alamos.

Recommendation 4-13

We recommend that the NAIRR specifically characterize the three levels, e.g., beginner, intermediate and advanced, if this is what is being described.

Recommendation 4-16

We recommend that NAIRR specify the architecture of edge computing resources and establish working relationships with edge computing resource suppliers to ensure continuity of resource availability over time.

Testbeds

Recommendation 4-19

IEEE-USA recommends requiring, not simply asking “when possible,” that this function be charged to the NAIRR staff who will be responsible for the implementation, management, and evolution of the testbed environment NAIRR chooses to utilize. This may guarantee a complete and widely accessible catalog, thus providing the needed consistency and avoiding duplication of efforts.

e. Privacy, civil rights, and civil liberties requirements. (Chapter 6)

Advanced analytics and artificial intelligence are powerful technologies that, along with their clear societal benefits, create new threats to privacy, equality, fairness, and transparency. Existing law does not yet protect sufficiently against these threats. The NAIRR does not mention ethics and appears to conflate privacy and security. IEEE-USA recommends adding clarifying language and reflecting on how NAIRR can manage the ethical use of these resources. The NAIRR could look to the recent work of both the Administration and Congress to ensure that AI systems adhere to accepted democratic standards of protections.

⁷ <https://alliancecan.ca/en/services/advanced-research-computing>

Additionally, NAIRR could refer to the work of scholars and practitioners from law, engineering, sociology, and statistics communities. For example, in a 2021 Ohio State University legal studies research paper, the authors interviewed corporate privacy managers, lawyers, and consultants, and surveyed a wide range of privacy managers to answer fundamental ethics questions about business data ethics management.⁸

Users of AI resources and tools must ensure that outcomes do not result in disparate treatment, disparate impacts, or other algorithmic harms and violations of democratic principles, and consequently undermine public confidence in and acceptance of AI. When AI systems are developed and deployed, objectives of accuracy and lack of algorithmic and other biases towards different groups can conflict. To mitigate these issues, it is imperative to ensure that researchers and students have access to established metrics and standards that will enable their operators to comply with standards that will enable their operators to comply with applicable legal and other standards for fairness, privacy, safety, and security. Transparency mechanisms for stakeholders that require third-party access to data in standardized, machine-readable formats are also needed.

Recommendation 6-1

NAIRR has noted that the protection of individuals' privacy, civil rights, and civil liberties are paramount to creating trustworthy AI and, therefore, the innovation, economic, and societal benefits that AI technologies and uses can deliver. Given the paramount importance of these protections under law, as well as under ethical systems, and given the devastatingly abundant evidence that government and private sector use of AI systems is falling short of the mark, Recommendation 6-1 should be strengthened to provide more substantive protections than are currently expressed in the accompanying text to take "efforts to ensure" transparency and "appropriate oversight" for NAIRR's operations, research, and governance.

Reaching beyond Recommendation 6-1, attorneys knowledgeable in privacy, civil rights, civil liberties, and consumer protection within the contexts of AI, specifically, and information technology, generally, should be integrated within NAIRR's operations, research, and governance teams.

IEEE-USA recommends incorporating references to laws and legal compliance. Specifically, transparent compliance with the Constitutional requirements for equality under the Fourteenth Amendment, the Civil Rights Act of 1964, and other civil rights laws should be specifically addressed and demonstrated. In addition, transparent compliance with the Freedom of Information Act (FOIA) should be highlighted, and, where exceptions under FOIA may apply, NAIRR should invoke and apply those exceptions only to the extent legally necessary to maximize transparency.

As to privacy, civil rights, and civil liberty protections within user agreements, the inclusion of terms is customary and generally assigned as an obligation to comply with the law. Mere agreements to comply, however, are inadequate. The law is not sufficiently evolved in its interpretation and application within AI contexts to make it clear to signatories how compliance is achieved. Therefore, we recommend that such user agreement terms provide specific guidance and include specific audit, reporting, and enforcement provisions.

Furthermore, user agreements should bar signatories from asserting or attempting to assert, including by registration or recordation, any property or other proprietary rights in resources made available by NAIRR thereunder. This bar should encompass a bar on attempts or assertions by those signatories to incorporate

⁸ Business Data Ethics: Emerging Trends in the Governance of Advanced Analytics and AI, *Ohio State Legal Studies Research Paper No. 628*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3828239. Accessed 27 June 2022.

those resources into their intellectual property assets, trade secrets, confidential information, or proprietary information, including as they may subsequently attempt to broadly define those terms.

Lastly, user agreements should incorporate compliance, audit, reporting, and enforcement provisions to protect data about individual human beings, irrespective of whether such data have been de-identified. Such privacy terms should not rely upon the now-outdated “reasonably linkable” standard articulated by the U.S. Federal Trade Commission in its 2012 Report to Congress.⁹

Recommendation 6-2

IEEE-USA recommends that ethics reviews under the recommended ethics review process be made iteratively over successive periods, for example, annually. Complete reports, findings, and corrective actions should be timely prepared and made publicly available.

In addition to an ethics review process, the NAIRR should establish a legal compliance review process. This legal review process should be carried out prior to resources being included within the system and then periodically thereafter. The legal review process should operate in tandem with the ethics review process. As a best practice, the earlier in the resource development process that legal and ethics experts can be engaged and integrated into the development teams, the sooner and more comprehensively complications can be successfully addressed.

As to third-party resources that may be targeted to be made available via the NAIRR, a thorough legal review should be carried out to document the provenance of those data and ensure that third parties are legally entitled to make those resources available to and via the NAIRR without breach of contract, privacy, or copyright infringement, and other legal violations, including moral rights that may apply.

As to “higher risk data,” embedded privacy protections alone may be inadequate to protect the privacy, civil rights, and civil liberties of the human individuals about whom the data relate. If these embedded protections were defective, inadequate, outdated, or defeated by cyberattacks, the individuals’ privacy, civil rights, and liberties protections would be potentially irretrievably harmed and unrecoverable. Therefore, the NAIRR should strongly consider a multi-layered protective approach to incorporate zero-trust models and homomorphic encryption, for example, in addition to these embedded protections.

The NAIRR should incorporate detailed mechanisms to track access to and use of these higher risk data.

The NAIRR has stated in the text accompanying Recommendation 6-2 that it may use institutional review boards (IRBs) for its research. The protective principles underlying IRBs should be adopted and as deemed appropriate, incorporated within the NAIRR’s recommended ethics review process. Under no circumstances, however, should ethics and IBR reviews be substituted for the legal compliance reviews and review process suggested, *supra*.

In the text accompanying Recommendation 6-2, the NAIRR expresses that, to the extent feasible, the outcomes of research that it enables should be regularly and over the long-term vetted to ensure that the subject research is not causing or contributing to violations of privacy or civil rights or infringements of civil liberties. The auditing, reporting, and enforcement mechanisms discussed in the commentary provided, *supra*, as to Recommendation 6-1 and user agreement terms will enable this essential vetting by the NAIRR and render that vetting feasible. We strongly recommend including those mechanisms.

⁹ Federal Trade Commission, *Protecting Consumer Privacy in an Era of Rapid Change*
<http://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf>. Accessed 30 June 2022.

Any research, however, that is carried out, without such user agreements in place must also be covered by vetting and protective mechanisms. Impact assessments such as those proposed in the 2022 Algorithmic Accountability Act¹⁰ or modeled in the Ethics and Algorithm Toolkit from The Johns Hopkins University Center for Government Excellence, for example, may be useful for the NAIRR to consider in this regard or more broadly.¹¹

Recommendation 6-3

We support Recommendation 6-3. Training, however, should be iterated on a quarterly basis as a best practice. All training should include a non-trivial assessment to confirm that trainees have a demonstrable and full understanding of their obligations, including their legal, fiduciary, and other duties of care to protect individuals' privacy and civil rights and their civil liberties.

We also recommend considering the establishment of NAIRR certifications that go beyond the expected scope of training. More rigorous certifications could more effectively and expeditiously recover and build greater public and market trust in AI systems and uses.

Recommendation 6-4

The NAIRR should ensure the completeness, currency, easy searchability, and free and public availability of the full listing and description of each of the planned inventory of provided data sets and dated details as to the history and provenance of each. As to retired and deprecated datasets, those should be maintained, within the inventory and its listing, while being designated as inactive.

¹⁰ See Algorithmic Accountability Act of 2022, S. 3572, § 4(a)(11)(b), 117th Cong., 2nd Sess. (introduced Feb. 3, 2022). A companion bill was introduced in the House of Representatives. See Algorithmic Accountability Act of 2022, H.R. 3580, 117th Cong., 2nd Sess. (introduced Feb. 3, 2022).

¹¹ Ethics & Algorithms Toolkit: A risk management framework for governments (and other people too!), <https://ethicstoolkit.ai/>, Accessed 29 June 2022.

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Internet2

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



June 30, 2022

The National Science Foundation
2415 Eisenhower Avenue
Alexandria, VA 22314
NAIRR-responses@nitrd.gov

Re: RFI Response: National AI Research Resource Interim Report

Dear Colleagues:

The University Corporation for Advanced Internet Development (d/b/a “Internet2”) appreciates the work of the Task Force and the opportunity to provide comments.¹ Internet2 is encouraged by the Task Force’s recognition of the important role that existing campus-, regional-, and national-scale resources play in supporting research activities as noted in Recommendation 2-3. As one of these national-scale resources, Internet2 is supportive of the interim report and takes this opportunity to provide additional comments in relation to topic letter b in the RFI relating to the establishment and sustainment of the NAIRR, including agency roles, resource ownership and administration, governance and oversight, resource allocation and sustainment, and performance indicators and metrics.

Recommendation 3-10: Access to NAIRR resources should be contingent on research project proposal review, be governed by clear use policies and user agreements, and be in compliance with relevant requirements for open sharing of research outputs.

The Interim Report recommends that “the NAIRR management entity should aim to implement standard legal agreements for users and resource providers, establishing common terms of use.” Further, the report states, “[s]uch legal agreements have the potential to substantially reduce the administrative burden that researchers and their institutions would otherwise face in establishing agreements with multiple resource providers on a case-by-case basis.”

Internet2 suggests that NAIRR should work to integrate with existing agreements and technical implementations in place at academic institutions and incorporate any specifically negotiated terms for into existing agreements. In many cases, the agreements already in place for cloud services are heavily negotiated to meet a complex set of institutional requirements, state laws and requirements, and federal laws. There also are a number of collaborative agreements for cloud services that are widely adopted across research and education institutions, including those negotiated as part of the Internet2 NET+ program.

During more than a decade of working with research and education institutions on collaborative cloud agreements, Internet2 has observed that the requirement to enter into new agreements for cloud services, in addition to those requirements that universities already maintain, creates procurement and technical overhead that slow speed to adoption and can potentially raise significant administrative, technical, and

¹ Internet2 submitted comments to the National Science Foundation on September 29, 2021, in response to the RFI on the National AI Research Resource. See: <https://www.ai.gov/rfi/2021/86-FR-39081/I2-NAIRR-RFI-2021.pdf>. In addition, Ana Hunsinger, Vice President of Community Engagement, participated in a panel discussion on *User Resources: Portal Interface and Educational Tools*, on October 25, 2021, elaborating on the recommendations offered by Internet2 in the September 2021 comments.



security challenges. Most major research institutions already have at least one, and likely more than one, implementation of public cloud infrastructure services in place. Thirty-two unique higher education institutions responded to the 2021 Cloud Forum Survey.² Of those, six institutions were single-cloud, with single-cloud defined as having no more than one cloud platform at five percent or more usage. The rest had made meaningful investments in two or more cloud platforms.

Internet2 has seen first-hand the challenge of tying a program to a specific contract or reseller and worked collaboratively to resolve that challenge. As an example, Internet2 has worked with institutions to streamline their access to the National Institutes of Health (NIH) STRIDES program in a manner that did not require: (1) agreeing to an additional contract; (2) creating a new business relationship with a channel partner (e.g., reseller); and/or (3) duplicating existing enterprise controls already managed by the institution. Working collaboratively with higher education institutions, the cloud service providers, channel partners, and NIH, Internet2 was able to implement this approach for both Amazon Web Services (AWS) and Google Cloud Platform.³

In a May 2021 letter to NIH, the NET+ AWS Service Advisory Board described the challenges caused by the requirement to enter into a new agreement to leverage STRIDES awards with AWS.⁴ Two excerpts from that letter are included below for the Task Force's consideration:

- (1) Our institutions have invested time and effort to integrate AWS into our technical, security, and business processes, enabling our researchers to efficiently use AWS in ways that integrate with established security policies and billing structures.
- (2) Use of a separate reseller would require deployment and maintenance of an entirely separate set of AWS security, networking, monitoring, and support infrastructure on our campuses, plus review, approval, and maintenance of a second AWS contract. This creates duplicate processes and technologies that rapidly become difficult to maintain, creating very real business and security risks.

In summary, the Task Force's goal of enabling pathways into AI research is best served by integrating into pathways that already exist in the context of cloud agreements. Many higher education institutions already have complex legal agreements and technical environments to enable public cloud resources for their users.

Respectfully submitted,

/s/ John S. Morabito
John S. Morabito
Vice President, External Relation

² See [2021 Cloud Forum Data Survey Results](#)

³ [Internet2 Announces Availability of Collaborative Agreement With AWS for NIH Award Recipients to Utilize the STRIDES Initiative - Internet2](#); and [Internet2 NET+ Google Cloud Platform Terms of Service Now Available for NIH Award Recipients to Utilize the STRIDES Initiative - Internet2](#)

⁴ [AWS NET+ Service Advisory Board - NIH Request](#); and [NET+ Amazon Web Services \(AWS\)](#)

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

SeedAI

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

SeedAI – NAIRR Interim Report RFI response

Please see the following comments on behalf of SeedAI. Let me know if you have any questions.

Thank you!

Austin

--

SeedAI thanks the National Science Foundation (NSF) and the Office of Science and Technology Policy (OSTP) for the opportunity to express our views on the National Artificial Intelligence Research Resource (NAIRR) Task Force interim report.

As an organization, we work with public/private partners to build community-centered AI ecosystems to support their interests and economic opportunity in the age of AI. We recently launched our AI Across America project in conjunction with the Congressional AI Caucus. Through this initiative, we're visiting a diverse array of communities to discover frameworks that work for them.

While it's early, we're finding broad agreement that the United States needs a NAIRR to account for the core of broadly-accessible AI resources and connective tissue that can serve as a foundation for innovation, investment, and partnership to bring AI in reach for people outside traditional tech hotspots and groups.

The Co-Chairs and the NAIRR Task Force have operated openly, intentionally, and in good faith towards a blueprint for the system. As a result, the framework expressed in the interim report is a solid foundation for the rest. We are fully supportive of the current direction, and look forward to its evolution into something Congress can concretely authorize and fund.

We recognize that the thorniest work begins now. The Task Force's framework for the NAIRR is necessarily grand in scope; further scoping and implementation will have to address core questions across AI research that extend to the field at large. The report identifies dozens of issues, regarding on data, compute, talent, safety, civil rights, bias, inclusivity, intellectual property, cybersecurity, public/private engagements – and that is only scratching the surface.

The effort is necessary and valuable; a fully realized NAIRR will be a cornerstone of U.S. competitiveness and global AI leadership, propelling domestic progress in AI. It will resource the genius of neglected communities across the country, and pave the way for a new wave of partnerships, discoveries, and economic growth.

Below are some preliminary thoughts based on our conversations to date. As we visit more communities as part of the AI Across America project in conjunction with the Congressional AI Caucus, we will expand and refine these suggestions for the Task Force's consideration.

- 1.
- 2.
3. Quickly explore questions around interoperability, and issue guidance for other related federal projects to ensure they can
4. easily integrate.
- 5.
- 6.
- 7.
8. Streamline the connection between the NAIRR and broader innovation investments, such as NSF Engines, state/local programs, and
9. private efforts.
- 10.
- 11.
- 12.
13. Launch the NAIRR with resources and challenge sets that will be immediately useful to the community and the function of the
14. NAIRR itself.
- 15.
- 16.
- 17.
18. Engage with diverse communities from the beginning, tailoring the approach and offering a clear value proposition.
- 19.
- 20.
- 21.
22. Create a streamlined mechanism for nonprofits to facilitate NAIRR access for startups, especially those underserved.
- 23.
- 24.
- 25.
26. Engage the community with contests and grants to research, develop, or apply technologies that support NAIRR operation.
- 27.
- 28.
- 29.
30. Prioritize sufficient staffing for NAIRR and NAIRR-supporting agencies, taking full advantage of the Intergovernmental Personnel
31. Act and other detailee programs.
- 32.

Thank you again for the opportunity to engage.

Sincerely,

Austin Carson
SeedAI | Founder & President

--



Austin Carson
SeedAI Founder & President
c: (202) 656-4210 | [Website](#)



Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

**Yonadav Shavit, Divyansh Kaushik, Zachary C. Lipton,
Samuel R. Bowman, Kira Goldner**

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

Dear Directors Parker and Parashar,

We congratulate you on the National Artificial Intelligence Research Resource (NAIRR) Task Force’s recently published Interim Report,¹ which we believe is a promising roadmap for substantially strengthening the American AI research ecosystem. We are a collection of academic AI researchers from universities across the country, united in our focus on techniques for building responsible and reliable AI. We write to address areas (a) and (c) of the Task Force’s report. The report highlights six Resource Elements and Capabilities that the NAIRR should include: general data, government data, compute resources, testbeds, user interfaces, and educational tools. We write to highlight a vital yet currently missing seventh resource needed to enable cutting-edge AI research, which we believe the NAIRR must provide: scalable interaction with humans, acquired through “data work” vendors. To ensure AI research moves towards building systems that are safe, accountable, and useful to humans, we must make frequent human interaction a ubiquitous element of AI training, deployment, and benchmarking/evaluation.

Summary:

- The National Artificial Intelligence Research Resource (NAIRR) should provide **funding and infrastructure** to researchers and small businesses who seek to utilize human interactions in their AI research, by negotiating centralized contracts with data work platforms.
- NAIRR should work with the National Institute for Standards and Technology (NIST) and the Department of Labor to promulgate **standards for ethical data work platforms**, and identify vendor options that comply with these standards.

Effectively incorporating human interaction and oversight is the next frontier of artificial intelligence. Despite the many breakthroughs achieved by machine learning (ML), a significant obstacle prevents its wider deployment in consequential applications: safety, accountability, and human compatibility. If we want to build ethical, safe, and accountable AI systems, we need to scalably incorporate human interactions into their development and deployment.

Data worker interactions play critical roles in AI research and development. Positive human evaluations can serve as part of an AI system’s objective. Humans need to monitor and audit AI behaviors, and developing tools to enable such monitoring requires constant interaction with data workers. Data workers can even provide ideal behavior demonstrations for the AI systems to learn from (known as “imitation learning”).² Researchers and practitioners at major companies are already widely adopting these practices, and they have contributed to several breakthroughs in human-compatibility of advanced AI systems.³ The Chinese AI R&D ecosystem similarly relies on large-scale usage of data workers in the development of cutting edge AI systems. Given the ways that these human-centered techniques may ameliorate or exacerbate the social, ethical, and safety challenges of AI systems, we do not want this research to be restricted to large companies.

¹ NAIRR Task Force. *Envisioning a National Artificial Intelligence Research Resource (NAIRR): Preliminary Findings and Recommendations*. May 2022

² Hussein, Ahmed, et al. "Imitation learning: A survey of learning methods." *ACM Computing Surveys (CSUR)* 50.2 (2017): 1-35.

³ Several major AI companies have publicly disclosed that they rely on repeated human feedback to make their AI systems behave more responsibly. These include Google’s dialogue system [LaMDA](#), OpenAI’s [InstructGPT](#) and [PALMS](#), and Anthropic’s [HHH projects](#).

What roadblocks prevent academic AI researchers and small businesses from utilizing data workers when developing AI systems, relative to their industry peers?

- First, procuring data worker labor while providing fair pay is expensive, especially for small teams who **cannot negotiate bulk rates**.
- Second, providing **ethical working conditions** to data workers is complicated, especially for projects that lack dedicated staff to oversee data workers. Further, academic researchers must further navigate time-intensive approvals from Institutional Review Boards.
- Third, the vast majority of AI researchers **do not possess the infrastructure** to integrate data workers effectively into research workflows, or to handle routine tasks like quality assurance. This difference in infrastructure is dramatic when compared with the infrastructure that exists in the rest of the AI field, e.g. for testbeds or machine learning software libraries.

At the same time, media reports⁴⁵ indicate pervasive poor working conditions for data workers, whose labor plays a key role in advancing the AI state of the art. Given the increasing prevalence⁶ of crowd-work as a source of income for everyday Americans and others across the world, it is important to ensure that data work creates good jobs. Additionally, as academics turn to the workforce in other countries for data labeling (presumably due to lower costs of annotation), each group's independently monitoring ethical standards can be difficult. Data workers' efforts determine AI systems' behavior; treating those workers ethically is the foundation of building ethical AI. While the tech industry is beginning to develop standards, the US government itself is best positioned to define standards that balance the interests of both companies and workers.⁷

The National AI Research Resource has a unique opportunity to *simultaneously* address both (1) the difficulties of access to large-scale human interactions in academic AI research *and* (2) to directly promote a more ethical data work ecosystem. By providing academics and small businesses with funding for data work, and conditioning that these funds/credits can only be spent on vendors that comply with ethical standards drafted by NIST, the government can support the emergence of **a market for ethical data work**. In time, these standards may spread well-beyond the NAIRR and come to serve as the industry gold standard.

Below, we list four recommendations that the National AI Research Resource Task Force should consider adopting to ensure American AI research can stay competitive on the new frontier of human-involved AI.

⁴Schmidt, Florian A. "Digital labour markets in the platform economy." *Mapping the Political Challenges of Crowd Work and Gig Work* 7 (2017): 2016.

⁵ "Facebook to Pay \$52m to Content Moderators over PTSD." *BBC News*, BBC, 12 May 2020, <https://www.bbc.com/news/technology-52642633>

⁶ Strozzi, C., and M. Cantarella. "Workers in the crowd: The labour market impact of the online platform economy." (2021).

⁷ "Responsible Sourcing of Data Enrichment Services." Partnership on AI, 27 May 2022, <https://partnershiponai.org/paper/responsible-sourcing-considerations/>.

Recommendation 1: The National AI Research Resource should engage in centralized procurement of “data worker” contracts (including pre-purchased “credits”) from data work platforms, and distribute the procured crowdworker-hours to researchers and small businesses.

The federal government can negotiate on behalf of academic researchers and small businesses to get better contracts for sourcing crowdworkers to participate as subjects and collaborators in human-feedback-based AI R&D. There are clear benefits to centralized procurement, including negotiating better rates, ensuring better labor conditions for workers, and reducing the administrative burden for technologists unused to collaborating with data workers. One simple contract structure would be to purchase a minimum guaranteed quantity of data worker hours, in exchange for locking in a favorable rate.

Any such contracts with data work vendors should include stipulations that:

- The vendors receiving these contracts must **comply with ethical standards** as outlined in “Recommendation 2” below.
- **Funds should not be pre-attached to vendors**, so that researchers and small businesses can shop around for whichever qualifying vendor best meets their needs.
- Many AI development projects require a fixed pool of human annotators who have been provided with some sort of dedicated training by the research team, or who otherwise possess unique skills or knowledge (e.g. dermatologists). Finding such workers may be impossible on crowdwork platforms, and the NAIRR should provide **alternative channels for researchers to identify these workers at their own discretion**.
- Academic researchers receiving these funds should **publicly release their collected human interaction datasets** in line with Section 4 of the White House Memo on “Increasing Access to the Results of Federally Funded Scientific Research”, unless otherwise prohibited.⁸

Recommendation 2: The National AI Research Resource should work with the National Institute for Standards and Technology (NIST) and the Department of Labor to establish standards for ethical AI data work.

A large US federal customer like the National AI Research Resource has the ability to set an example for industry, and the data work vendor industry is in dire need of clear standards to follow. In drafting standards that prospective contracting vendors must meet, the National AI Research Resource could collaborate with the National Institute of Standards and Technology (NIST) and the Department of Labor. There are several specific areas where formal standards would substantially improve both data quality and working conditions:

- Defining tiers of hazardousness for data work, and appropriate platform precautions for each tier (e.g. involving labeling of violent videos).
- Outlining best practices on the disclosure of the composition and demographics of data workers, especially when their contributions substantially shape the oversight of an AI system.
- Defining tiers of data quality assurance, including rates of fraud or abuse.

⁸ [John P. Holdren, "Increasing Access to the Results of Federally Funded Scientific Research", February 22, 2013](#)

- Defining best practices on processes for recourse and arbitration between data workers, platforms, and data customers (like AI researchers or small businesses).

We further recommend that NIST or another federal entity accredit and publicize the list of data work vendors that meet these standards. This could serve as a gold standard accreditation for private companies who wish to ensure the ethical compliance of their AI data work.

There are other data work procurement practices that are difficult to standardize, but in which the NAIRR can nevertheless lead by example. In particular, the NAIRR can prioritize providing fair pay to data workers, and place special emphasis on pressuring data vendors to address human rights concerns.

Recommendation 3: The National AI Research Resource should work with the National Science Foundation (NSF), Department of Health and Human Services (HHS), and the Food and Drug Administration (FDA) to create a set of umbrella IRB templates/pre-approvals for research projects that use human feedback to follow one of a set of pre-established categories.

Institutional Review Board (IRB) reviews of human subject research are necessary for all academic human feedback if researchers use, study, or analyze information about living individuals. However, the complexity of IRB approval imposes a disincentive for unfamiliar researchers⁹, which limits adoption among researchers developing consequential AI systems which would benefit from human oversight. The NAIRR should work with the NSF, HHS and FDA to create a set of standard IRB templates for canonical machine learning data worker tasks, provided that researchers comply with pre-defined requirements. These agencies could also promulgate clearer guidelines on which canonical ML research tasks are fully exempt.

Recommendation 4: As part of its “User Interface” resources, the National AI Research Resource should promote an ecosystem of tools that help academics and small businesses to automate the logistics of managing human feedback experiments.

A major component of incorporating human interaction into AI systems is building the logistical infrastructure for handling the contributions of data workers. Many companies using data workers re-develop custom tools in-house, but individual researchers and small businesses cannot afford this overhead. A few examples include tools for managing fraud and abuse detection, and custom user interfaces for common data work tasks like annotating video or chatbot interactions.

The National AI Research Resource can effectively mitigate this burden by supporting the development of an open-source ecosystem of tools for data work infrastructure, and allowing these tools to directly integrate with the NAIRR’s “User Interface”. The NAIRR should not build the tools themselves, but should support external open-source tool developers as was done with the successful NASA TOPS program¹⁰, including by standardizing application programming interfaces (APIs) for integration with data work vendors.

⁹ Kaushik, D., Lipton, Z. C., & London, A. J. (2022). Resolving the Human Subjects Status of Machine Learning's Crowdworkers. *arXiv preprint arXiv:2206.04039*. <https://arxiv.org/abs/2206.04039>

¹⁰ [Transform to Open Science \(TOPS\)](#)

We believe that providing US researchers and small businesses with human feedback will substantially accelerate the development of safe, accountable, and useful AI systems. We look forward to the NAIRR Task Force's response, and are eager to assist further in ensuring the continued strength of American AI research. If you have any questions, please reach out to Yonadav Shavit (yonadav@plaintextgroup.com) or Divyansh Kaushik (dkaushik@fas.org).

Signed,

Yonadav Shavit,
Doctoral Candidate, Harvard John A. Paulson School of Engineering and Applied Sciences
and
Associate, Schmidt Futures

Divyansh Kaushik,
Doctoral Candidate, Carnegie Mellon University School of Computer Science
and
Science and Technology Policy Fellow, Federation of American Scientists

Zachary C. Lipton,
Assistant Professor of Machine Learning and Operations Research
Carnegie Mellon University

Samuel R. Bowman
Assistant Professor of Data Science, Linguistics, and Computer Science
New York University
and
Member of Technical Staff (sabbatical visitor)
Anthropic

Kira Goldner
Shibulal Family Career Development Assistant Professor of Computing & Data Sciences
Boston University

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Matt Sheehan, Andrew Critch, Krystal Jackson, Jacob Feldgoise

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

1779 Massachusetts Avenue, NW

Washington, DC 20036

P +1 202 483 7600 F +1 202 483 1840

CarnegieEndowment.org

June 30, 2022

Compute Accounting

Response to [RFI](#) on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force

The authors include an AI research scientist as well as policy researchers who specialize in AI policy. This comment pertains to topic letter “b” in the RFI—specifically the subsections on “resource allocation and sustainment” and “performance indicators and metrics.”

In order to implement Recommendation 3-17¹ of the [interim report](#), the NAIRR management agency should develop and implement a system for “compute accounting,” standardized methods to track and audit the use of computational resources, analogous to Generally Accepted Accounting Practices (GAAP) for financial resources. As a federated mix of computational resources, the NAIRR will need a standardized system to track and audit compute usage across each component resource—ensuring that NAIRR resources are equitably distributed and are not misused.

Technology companies today already build tools to internally track compute resource usage. For example, in algorithmic stock trading, it is not uncommon to maintain a company-wide dashboard displaying how much compute is being used by which algorithms and for what purposes. Cloud computing platforms such as [Google Cloud](#) and [Amazon Web Services](#) offer similar services, but

¹ “Data should be gathered to support four levels of performance indicators [including] ... (2) measures of resource usage/activities, including user diversity”



exact methods vary from company to company, and there is currently no industry-wide standard for compute accounting.

A compute accounting system implemented by NAIRR should incorporate two core features: measurement of floating point operations and scale sensitivity. Floating point operations (FLOPs) are the most widely-used metric for calculating compute expenditures because they are easily convertible across resource types and applications, at least within an order of magnitude difference in significance. The NAIRR’s compute accounting systems should similarly use FLOPs — or approximately-fungible equivalents, such as multiply-accumulate operations (MACs) — when estimating the computational expenditure for a given user or project. With FLOPs as a standard unit of measure, NAIRR’s compute accounting system could then implement scale sensitivity in its documentation requirements. Just as financial accounting practices have “materiality” thresholds, NAIRR’s compute accounting system should also be scale-sensitive—asking for greater documentation and transparency for large compute expenditures than for small ones.

If the NAIRR management agency knows how much compute was expended to train each ML model, it can estimate the risk of misuse and assess whether compute is equitably distributed. Larger models, which require more computational resources, generally carry higher risks. The NAIRR management agency could conduct internal audits of the largest (highest risk) projects, to ensure those allocated resources are not being misused. Furthermore, paired with demographic information on NAIRR users, compute accounting data could be used to calculate the share of NAIRR resources that is supporting researchers from traditionally underserved communities.

While compute accounting will mostly help NAIRR internally monitor resource usage, NAIRR can also serve as an example for a standardized, replicable approach to compute accounting. As a leading public compute resource, the NAIRR would be in a good position to develop a practical and easily accessible method of compute accounting that could in turn be adopted across industry and academia. Doing so would scale up the equity and safety benefits generated by compute accounting, and set a positive example for other institutions.



We encourage the NAIRR Task Force to incorporate compute accounting in its final report by adding the following to Recommendation 3-17:

1. The NAIRR management agency should work closely with NIST and industry leaders to develop:
 - a. Standardized “compute accounting” methods that can be used to track resource usage across multiple resource types, professionally analogous to Generally Accepted Accounting Practices (GAAP) for financial accounting;
 - b. Standardized auditing procedures to screen for and detect dangerous applications of computational resources.
2. For large expenditures of compute, the NAIRR management agency should:
 - a. Require greater documentation and transparency,
 - b. Conduct regular audits.

Thank you very much for your attention, and please don't hesitate to reach out if you would like to further discuss this proposal.

Sincerely,

Matt Sheehan

Fellow, Carnegie Endowment for International Peace

Andrew Critch

Research Scientist, UC Berkeley Department of Engineering and Computer Sciences

Krystal Jackson

M.S. Information Security Policy & Management
Carnegie Mellon University

Jacob Feldgoise

Junior Fellow, Carnegie Endowment for International Peace

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Software & Information Industry Association (SIIA)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



RFI Response: National AI Research Resource Interim Report

Submitted by the Software & Information Industry Association to the Office of Science and Technology Policy and the National Science Foundation

On behalf of the Software & Information Industry Association (SIIA), we appreciate the opportunity to provide feedback on the interim report of the National Artificial Intelligence Research Resource (NAIRR) Task Force (the Task Force).

I. Preliminary remarks

SIIA, a non-profit organization, is the principal trade association for the software and digital information industries worldwide. Our members include over 450 companies reflecting the diversity of the information landscape, from creation to dissemination to productive and responsible use. They include digital content providers and users in academic publishing, education technology, and financial information, along with creators of software and platforms used by millions worldwide, and companies specializing in data analytics and information services. Our members support policies that foster innovation and a healthy digital ecosystem, including consumer privacy protections, responsible and ethical AI, and diversity, equity, and inclusion (DEI) initiatives.

We congratulate the Task Force for its excellent report. The report demonstrates robust engagement with the challenges of expanding access and pathways to ensure a more diverse, equitable, inclusive, and accessible AI R&D ecosystem in the United States. We view the establishment of a NAIRR as a critical means to achieve these goals.

We endorse the objectives and vision as the Task Force has presented them in Chapters 1 and 2. Strengthening the U.S. innovation ecosystem to realize the potential of AI applications to advance “science, economic growth, national security, and the ability to meet pressing societal challenges” while protecting privacy, civil rights, and civil liberties. Our collective ability to unpack this potential most fully will depend on our ability to educate and foster talent across socio-economic lines, particularly, as the report notes, among “traditionally underrepresented groups in AI R&D” (Rep. at 2-1), and to provide access to data and compute resources. This approach will help to advance DEI goals with respect to AI expertise and will provide a stronger foundation for ensuring that values-based assessments are built into the datasets and algorithms that drive AI applications.

We recognize that the Task Force intends the report as an outline for the NAIRR project. As the Task Force works to develop implementation steps to realize the NAIRR, we provide the following comments in the spirit of assisting in that process and address only a handful of the recommendations.

II. Establishment and sustainment of the NAIRR (Chapter 3)

The report demonstrates careful consideration of alternative options for structuring the NAIRR. Each of the alternatives presented in [Recommendation 3-3](#) has benefits and downsides. Considering these alternatives in light of the Task Force’s vision and objectives for the NAIRR, we recommend that the Task Force propose to establish the NAIRR within an existing FFRDC. We believe this approach would provide the NAIRR with the strongest path to achieve its objectives in a manner that promotes diversity and democratization of AI R&D in the United States. It will speed the process from legislation to execution by leveraging the FFRDC’s existing back-office infrastructure and expertise in managing complex government and public-private projects.

Locating the NAIRR in the Federal government has appeal although will present extraordinary challenges with respect to basic organizational requirements. These include appropriations restrictions, limited ability to obtain resources and funding from non-governmental sources, and restrictive hiring authorities. These challenges are likely to delay creation of the NAIRR and generate ongoing complexities in execution.

A university-based approach also has appeal. However, as the report notes, there is a growing divide in AI resources concentrated in large private-sector firms, well-resourced universities, and national laboratories.” (Report at 1-1.) The approach taken by the National Science Foundation (NSF) AI Institutes program can serve as a model to democratize access to AI R&D resources, although the decentralized nature of the program will present other challenges in executing the NAIRR vision.

Assuming the Task Force will recommend establishment of the NAIRR outside the Federal government, we would encourage the Task Force to provide additional guidance on the role of Federal agencies in the establishment and ongoing operations of the NAIRR. [Recommendations 3-1 and 3-2](#) call for involvement by multiple Federal agencies, including the Department of Energy, the NSF, the National Institutes of Health, and the National Institute of Standards and Technologies (NIST). This “federated approach” is critical to facilitate oversight and guidance, expertise, funding, access to government data, and other resource needs.

We believe attention to the role of the Federal government is essential because the Federal government has an unmatched ability to catalyze diverse actors, including research institutions, private industry, civil society organizations, and state and local governments. This convening power is what distinguishes the NAIRR from other efforts to advance and democratize AI R&D.

While the Task Force recommendations make clear that Federal agencies will have ongoing roles in providing access to government data and agency expertise, we encourage the Task Force’s implementation plan to include recommendations on the following:

- Identifying a lead agency or office. The success of the NAIRR in the short term will require dedicated guidance and participation of the Federal government. Though a federated approach makes sense from the perspective of resourcing, coordination will be essential to assisting NAIRR in coordinating among different stakeholders and marshalling the resources of the Federal government. An entity such as the Office of Science and Technology Policy (OSTP) or the National Science Foundation (NSF) has requisite experience to take on this role.
- Anticipated role of Federal agencies in supporting the NAIRR at establishment and in an ongoing manner. Members of the Task Force have the expertise and experience to provide a concrete vision for how the Federal government will engage with NAIRR at inception and over time. Providing guidance to implementers with further detail on agencies' anticipated financial and resource contributions will help to guide Congress, the Executive Branch, and the NAIRR management entity in advancing the NAIRR proposal.
- Anticipated need for specific Federal roles to support NAIRR. The NAIRR should seek to leverage government expertise in critical areas. For example, we would recommend that NIST lead the effort to provide standards for assessing the quality of data pools contributed to and created by the NAIRR. The NAIRR should leverage NIST's expertise in developing standards for test, evaluation, verification, and validation procedures and building a risk-management framework for responsible AI.

In addition, [Recommendation 3-1](#) contemplates that Congress will appropriate funding to individual Federal agencies that will support the efforts of NAIRR and that NAIRR management will explore additional, presumably non-governmental, revenue sources. We support the approach to funding NAIRR through multiple sources. As funding (and other resource support) will be fundamental to success, we would urge the Task Force to consider funding needs over a five-to-ten-year period with recommendations about the level of funding that may be required from Congress and from private sources.

Consistent with our remarks regarding the unique convening power of the Federal government, the Task Force should consider what ongoing role Congress may have in sustaining the NAIRR infrastructure and providing an ongoing appropriations source – either directly or, as currently proposed, through individual Federal agencies. [Recommendation 3-16](#) would require dissemination of reports to the public, Congress, and supporting Federal agencies. Beyond this, the report does not contemplate a role for Congress. While it may be left to the management entity to determine funding needs and sources, on the assumption that regular appropriations from Congress will be needed, we

encourage the Task Force to consider what ongoing relationship the NAIRR will have with Congress in terms of oversight and potentially direct appropriations.¹

Recommendations 3-13 and 3-14 address contributions from the private sector, which will be essential to ensure the success of NAIRR in resource intensive areas (such as data and compute, covered in Chapter 4) regardless of congressional appropriations. We support involvement of the private sector and note that several firms, during the initial RFI period, indicated willingness to support the NAIRR through different types of contributions. We encourage the Task Force to provide further guidance about anticipated private sector financial and in-kind contributions, including anticipated needs and recommendations to address potential conflicts while allowing NAIRR to leverage private sector resources to support talent, data, and compute necessary for the NAIRR program.

We strongly endorse Recommendation 3-7's call “to build a DEIA focus into the system and operational plan from the beginning, rather than as an afterthought.” We support elevating and elaborating on this recommendation as it is critical to the success of the NAIRR in achieving its overall objectives. The Task Force may consider moving this recommendation into Chapter 2. As the Task Force develops its implementation plan, we urge it to consider recommendations specific to cultivating talent from underrepresented groups including from minority serving institutions.

III. NAIRR resource elements and capabilities (Chapter 4)

Our feedback on Chapter 4 focuses on the findings and recommendations with respect to data. Access to robust, reliable, and trustworthy data is a key impediment to the democratization and diversification of AI innovation and to the quality of AI innovation. Developing robust datasets that meet the standards for responsible AI and minimize privacy concerns is extremely costly for most researchers, state and local government agencies, and companies. The alternative of relying on poor quality data increases the likelihood of unintentional bias and faulty predictions. Datasets that do not comport with standards of accuracy, reliability, trustworthiness, and bias present significant societal risk.²

We strongly endorse the Task Force’s findings on data. As Findings 4-1 through 4-3 accurately claim, the curation and aggregation of robust, high-quality datasets is one of the leading challenges that

¹ Congress has in the past created non-governmental entities and has funded them in different ways. The Corporation for Public Broadcasting (created by the Public Broadcasting Act of 1967), for example, continues to receive Federal appropriations, while the National Constitution Center (created by the Constitution Heritage Act of 1988) received “seed” funding and now relies exclusively on philanthropic support, ticket sales, and membership.

² See, e.g., Joshua New, “AI Needs Better Data, Not Just More Data,” Center for Data Innovation (Mar. 20 2019), <https://datainnovation.org/2019/03/ai-needs-better-data-not-just-more-data/>; Tasha Austin, et al., “Trustworthy Open Data for Trustworthy AI,” Deloitte Insights (Dec. 10, 2021), <https://www2.deloitte.com/us/en/insights/industry/public-sector/open-data-ai-explainable-trustworthy.html>.

AI researchers and experts face when conducting their work. Differences in data labeling and data curation hinder the widespread adoption and deployment of AI in a variety of fields. In tandem with data challenges, many AI experts (particularly those in underserved and underrepresented areas) struggle with the acquisition of necessary computational resources. Finding 4-8 aptly describes a common occurrence in which AI experts may be hindered in their work without access to sufficient compute resources. Within a more democratized ecosystem that the NAIRR will provide, AI experts will be able to surmount these challenges, as resources—not skill—are the greatest limiting factor in the national AI R&D environment.

By their very nature, AI and ML technologies benefit from access to vast datasets. Whether through an independent aggregation of a large dataset or a multilateral conglomeration, AI stands to gain massively from diverse data. Independently aggregating data, however, can be extremely costly and difficult to accomplish. The alternative, while simpler with regards to the actual gathering of data, poses significant challenges that could be addressed by the NAIRR. Commonly, multilateral data aggregation suffers from issues such as the storage of data, differences in data labeling, and a lack of high quality, specialized data all of which originate from the decentralized nature of this approach.³ The NAIRR would aid in surmounting this problem by providing a central entity in which contributors could store and share data to facilitate a cooperative effort on research fronts.⁴ Furthermore, with a more unified structure, the NAIRR would offer the opportunity to present labeling standards to which data contributions must adhere, resolving the issue of heterogeneity. With contributing entities able to focus on specific data of their choosing rather than being concerned with quantity, this specialization could increase the overall quality of the NAIRR’s stored data. Experts in the field of AI allege that in recent years, enormous investments have been diverted away from AI R&D to other financial ventures.⁵ The proposed establishment of the NAIRR would thus aid in reinvigorating AI research and overcoming present roadblocks.

The research opportunities that the establishment of the NAIRR poses are significant. The opportunity for contributors to focus solely on their area of data expertise will yield greater quality, more reliable data. This refining of data presents an excellent opportunity for researchers and companies alike to conduct their own research on a large, robust dataset. Within the medical field, for

³ Sara Brown, “Why it’s time for ‘data-centric artificial intelligence,’” MIT Sloan (June 7, 2022), <https://mitsloan.mit.edu/ideas-made-to-matter/why-its-time-data-centric-artificial-intelligence>.

⁴ Connor Wright, “Our Top-5 takeaways from our meetup ‘Protecting the Ecosystem: AI, Data and Algorithms,’” Montréal AI Ethics Institute (Sept. 20, 2021), <https://montrealaiethics.ai/our-top-5-takeaways-from-our-meetup-protecting-the-ecosystem-ai-data-and-algorithms/>.

⁵ RE•WORK, “Experts Predict The Next Roadblocks in AI” (Aug. 20, 2020) <https://blog.re-work.co/experts-explain-the-next-roadblocks-in-ai/>.

instance, there presently is insufficient usable medical data for research.⁶ Some of the largest datasets available to medical experts are strewn across a mélange of “national government-sponsored studies, insurance claims, large clinical trials, cohort studies, and individual institutional registries.”⁴ To centralize the needed data in the NAIRR, actualization of Recommendation 3-13 would be an excellent mechanism to do so by linking an entity’s access to NAIRR resources to its data contribution levels. This approach would augment data available through the NAIRR while also providing much-needed computational resources. By amassing data not only across medical fields but also all disciplines while simultaneously encouraging that quality be upheld, the NAIRR presents a unique opportunity for large amounts of reliable data to be available, a key component in the democratization of artificial intelligence.

One of the NAIRR’s most effective tools in ensuring high quality data is contained within Recommendation 4-6. The provided technical infrastructure and support staff are integral to the NAIRR’s success, as they serve to educate and cultivate “community-driven standards and improvements to data quality as are determined by the relevant domains.” Taking the medical field as an example once more, there is no standard model or procedure by which medical experts gather data.⁴ To successfully deploy and scale AI research on a national level, leadership is required to standardize disease diagnoses (data categorization and labeling) in order to establish “ground truths” or “gold standards” for classification.⁷ These “ground truths” are of the utmost importance when running supervised AI models which are trained to predict and classify based on said truths.⁵ The NAIRR’s permanent technical support staff would be able to standardize data practices within appropriate domains, thus negating the need for post hoc data curation.

This robust, high-quality data that is aggregated within the NAIRR can then be analyzed and used by others. It is this form of multilateral data collaboration that can empower professionals to perform research previously found to be extremely costly. By democratizing and opening access to this data, the NAIRR could enormously expand the number and range of studies and research conducted.⁵ AI models themselves benefit from having a plethora of data sources, and a plethora of researchers would benefit from these new AI possibilities. The outline to incentivize data contributions as a collective (as alluded to in Recommendations 3-14 and 4-1) provides the NAIRR the chance to overcome a sort of collective action problem and to have the widest possible impact on the AI community.

Democratizing AI R&D has economic impact as well. The inherent nature of data is nonrival, and thus benefits can be derived from data aggregation at a large scale. Some experts at the Joint Research Centre (JRC) in the European Commission allege that from a societal perspective, “it may therefore be

⁶ Kobayashi, Y., et al., “How will ‘democratization of artificial intelligence’ change the future of radiologists?,” *Japanese Journal of Radiology* 37, 9–14 (2019), <https://doi.org/10.1007/s11604-018-0,793-5>.

⁷ Wang, Sophia Y et al. “Big data requirements for artificial intelligence,” *Current opinion in ophthalmology* vol. 31,5 (2020): 318-323, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8164167/>.

better to share data as widely as possible” as could be done through the NAIRR.⁸ Furthermore, the economic nature of data appears to align with economies of scope and scale. Those of scope focus on aggregation of data across a “variety of situations and observations” into a single dataset, while those of scale focus on large, in-depth datasets.⁶ The combination of these two (breadth and depth) are captured within the structure of the NAIRR in its appeal to varying disciplines and its dedication to quality.

We encourage the Task Force’s implementation plan to build out Recommendations 4-1 through 4-5 with concrete methods for the NAIRR to obtain and develop large, high-quality, and privacy-protected datasets. Despite the richness of government data and the thorough recommendations to make government data available through the NAIRR (see Recommendations 4-10 and 4-11), there remains a significant amount of data that is not within government control. The NAIRR can serve a critical function by gathering and making accessible data and incentivizing the creation of new datasets.

Specifically, we recommend that the Task Force include methods for leveraging private sector contributions, creating new public-private initiatives to gather and curate data, and leveraging Federal government expertise to ensure that datasets made available to researchers through the NAIRR meet high-quality standards. Examples of such methods include:

- Leveraging private sector resources to generate large synthetic data pools. Synthetic datasets can enable algorithms to run on data that reflect, rather than rely on, real-world data. This approach would allow for the creation of a robust data lake that can be vetted to ensure accuracy, reliability, fairness, and so on. Moreover, it would not present privacy and individual rights concerns that may arise from the collection, retention, sharing, and use of datasets that are built directly from personal information. We understand there is interest in the private sector to work with the government on this sort of initiative.
- Incentivizing private sector companies to provide unique data in a non-proprietary form. Many potential AI applications rely on proprietary data that private sector entities are understandably reluctant to make available. The NAIRR should explore methods to incentivize collection of such data on a voluntary basis and leverage appropriate privacy enhancing technologies (PETs) to ensure protection of proprietary information.
- Fostering the creation of large open datasets of personal information collected through enhanced notice and consent procedures. Personal information remains critical to many potential AI applications yet the collection and use of such information raise privacy concerns. Pilot projects to gather new forms of data (such as voice samples) from

⁸ Martens, Bertin, “The Importance of Data Access Regimes for Artificial Intelligence and Machine Learning,” JRC Digital Economy Working Paper 2018-09 (Dec. 2018), <https://ssrn.com/abstract=3357652>.



individuals who have received notice and consented would help to avoid these concerns.

IV. Privacy, civil rights, and civil liberties (Chapter 6)

We endorse the finding and recommendations in Chapter 6 of the report. We offer the following suggestions for consideration by the Task Force in developing the implementation plan.

First, we encourage the Task Force to consider proposing a framework for assessing privacy, civil rights, and civil liberties issues with respect to data and algorithms. Recommendation 6-2 nicely outlines an ethics vetting process. While there is wide agreement on the need for ethics assessments, there is variance on what benchmarks or criteria should be used to guide evaluation, particularly with respect to civil rights and civil liberties concerns. We recognize that the Task Force recommends that the NAIRR management entity have responsibility for developing these criteria. Given the extraordinary expertise and experience on the Task Force, this task would benefit from additional guidance from Task Force members.

Second, with respect to the vetting process outlined in Recommendation 6-2, we encourage the Task Force to examine what role the Federal government and other actors should have in ensuring that the data used within the NAIRR framework meet high-quality standards for reliability, trustworthiness, and bias. We encourage the Task Force to explore ways to incorporate NIST's standards and expertise into a vetting process.

V. Ideas for developing a roadmap to establish and build out the NAIRR in a phased approach, and appropriate milestones for implementing the NAIRR.

As the Task Force develops its implementation plan, we recommend preparing draft legislation and, as appropriate, text of proposed executive orders. We have found that providing lawmakers and policymakers with draft text is generally welcome and helps to facilitate the path from concept to realization.

* * *

Thank you for the opportunity to provide feedback on the interim report of the Task Force. We would be pleased to discuss any of these issues in further detail. Please direct any inquiries to **Paul Lekas**, SIIA Senior Vice President for Global Public Policy (plekas@siaa.net).

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Stanford Institute for Human-Centered Artificial Intelligence (HAI)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



Response to Notice of Request for Information (RFI) on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force
Stanford Institute for Human-Centered Artificial Intelligence

On behalf of the Stanford Institute for Human-Centered Artificial Intelligence (HAI), we are pleased to see that a large majority of recommendations in the interim report aligned very closely with our white paper, “[Building a National AI Research Resource: A Blueprint for the National Research Cloud](#).” We concur with the recommendations on the selection of management entity (3-4), tiered model for access to and storage of federal agency datasets (4-9, 5-1, 5-2), leveraging existing federal data sharing plans to facilitate access to datasets (4-11), and ensuring that NAIRR resources are allocated to specifically support research on AI trustworthiness (6-4).

We write to provide some additional feedback on specific recommendations for consideration in response to the Request for Information (RFI) by the National Science Foundation and the White House Office of Science and Technology on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force.

* * * *

Recommendation 2-6: While introducing commercial access to NAIRR for researchers at small businesses and private companies may very well benefit national AI innovation, we emphasize that the area of most acute need is for scientific research, particularly in higher education, and that an industry-oriented NAIRR poses substantial challenges. We recommend the Task Force limit access to NAIRR to researchers at U.S. higher education institutions during the first three years of a pilot run.

First, it is unclear how including small businesses and private companies meets the strategic objective of NAIRR. While we concur with the interim report that the mission of NAIRR is to expand and democratize access to AI R&D resources across the United States, it is important to provide ample infrastructure for basic scientific research which is a substantially under-resourced area in comparison to commercial research. As we noted in Chapter 1 of our white paper, the commercial sector is not part of the U.S. AI innovation system that is currently facing the structural challenges of lacking access to compute and data resources as well as the loss of AI talent.¹ Second, establishing NAIRR will be a complex and resource-intensive process. Introducing commercial access to NAIRR for small businesses and private companies that receive federal grants at launch may introduce a variety of regulatory and logistical challenges in the short term, further complicating and delaying the launch. For example, there are 31.7 million small businesses in the United States and federal agencies distribute on average 5,000-7,000

¹ Daniel E. Ho et al., *Building a National AI Research Resource: A Blueprint for the National Research Cloud* (Stanford Institute for Human-Centered Artificial Intelligence, Stanford University, October 2021), https://hai.stanford.edu/sites/default/files/2022-01/HAI_NRCR_v17.pdf.

awards per year via the Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) programs.² Opening up NAIRR access to those startups and other private sector researchers doing research that is in the “public interest” raises a wide range of boundary questions that NAIRR may be ill-equipped to adjudicate. Alternatively, the Task Force could consider extending subsidized loans to small businesses or private companies via other federal agencies outside NAIRR’s jurisdiction for purchasing computing resources to advance AI R&D.

Expanding NAIRR access to nonprofit organizations, federal agencies, or federally funded research and development centers (FFRDCs), as the interim report recommended, may be a more reasonable consideration and closer to the core of NAIRR’s mission. The Task Force could consider focusing on academic researchers as a starting point as it illuminates some of the main operational considerations for NAIRR access and adopts a broader access model in the long term.

Recommendation 4-12: We recommend the Task Force adopt a dual investment strategy with regard to computing infrastructure by developing programs for expanding access to existing commercial cloud services and building a high-performance computing (HPC) infrastructure to provide publicly owned resources. In the short run, scaling up cloud credit programs, using commercial cloud services (similar to NSF’s CloudBank program), provides numerous efficiency advantages. In the long run, our research shows that it is more cost-effective to own infrastructure when computing demand is close to continuous.³

Recommendation 3-10 & 6-2: For the proposal review and ethics review of researchers requesting NAIRR access, we recommend the Task Force adopt a tiered model. Researchers should gain access to base-level compute and data access by default and then apply through a streamlined process to gain access at resources beyond the base level on a project-specific base. A case-by-case, manual review for every single request for resource access, whether data or compute, would be an onerous process that balloon administrative overhead. Additionally, when researchers are simply applying for access, the research may be at an early stage without much to review.

For the ethics review specifically, researchers requesting beyond base-level compute should also be required to submit ethics impact statements with research proposals as part of the application. Existing mechanisms commonly used to assess academic research involving human subjects, such as institutional review boards, are ill-equipped to examine AI-related research as the research may not involve human subjects or rely on existing, publicly available data (not collected by the proposers) about people. In the meantime, the Task Force should consider establishing an ex post process to handle complaints about unethical research conduct or outputs.

* * * *

² “SBIR/STTR award data,” <https://www.sbir.gov/sbirsearch/award/all/>.

³ Preston Smith et al., “Community Clusters or the Cloud: Continuing Cost Assessment of On-Premises and Cloud HPC in Higher Education,” *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines* (July 2019): 1-4, <https://doi.org/10.1145/3332186.3333155>.

As lead authors, we proudly submit this response on behalf of our colleagues and the Stanford Institute for Human-Centered Artificial Intelligence (HAI).

Daniel E. Ho, J.D., Ph.D.
William Benjamin Scott and Luna M. Scott
Professor of Law, Stanford University;
Faculty Associate Director, Stanford Institute
for Human-Centered Artificial Intelligence
(HAI)

Jennifer King, Ph.D.
Privacy and Data Policy Fellow, Stanford
Institute for Human-Centered Artificial
Intelligence (HAI)

Russell C. Wald
Director of Policy, Stanford Institute for
Human-Centered Artificial Intelligence (HAI)

Daniel Zhang
Policy Research Manager, Stanford Institute
for Human-Centered Artificial Intelligence
(HAI)

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

The MITRE Corporation

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

©2022 The MITRE Corporation. ALL RIGHTS RESERVED. Approved for public release. Distribution unlimited. Case Number 21-01760-25



Response of The MITRE Corporation to the OSTP RFI on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force

June 30, 2022

For additional information about this response, please contact:

Duane Blackburn
Center for Data-Driven Policy
The MITRE Corporation
7596 Colshire Drive
McLean, VA 22102-7539

policy@mitre.org
(434) 964-5023

About MITRE

MITRE is a not-for-profit company that works in the public interest to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation. We operate six federally funded research and development centers (FFRDCs) that are leveraged by numerous federal agencies, participate in public-private partnerships (PPPs) across national security and civilian agency missions, and maintain an independent technology research program in areas such as artificial intelligence, intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience. MITRE's 9,000-plus employees solve problems for a safer world, with scientific integrity as our foundation. We are prohibited from lobbying, do not develop or sell products, have no owners or shareholders, and do not compete with industry. Our multidisciplinary teams (including engineers, scientists, data analysts, organizational change specialists, policy professionals, and more) are thus free to dig into problems from all angles, with no political or commercial pressures to influence our decision-making, technical findings, or policy recommendations.

Over the decades, MITRE has established and supported dozens of interdisciplinary, systems-level partnerships—such as collaboratives, consortia, and specialized PPPs—to bring whole-of-nation focus to achieving national priorities. Examples include the Aviation Safety Information Analysis and Sharing (ASIAS), the Medical Device Information Analysis and Sharing (MDIAS), and the Partnership for Analytics Research in Traffic Safety (PARTS). Because of our close relationships with federal agencies and our prohibition on competing with industry, we often serve these partnerships as a convener, building relationships among historically siloed groups, and as an independent steward of partners' proprietary/sensitive data—given that all parties can trust us to act in a conflict-free manner and focus solely on achieving national public interest objectives. Within ADAS, for example, MITRE supported the partnership by managing, safeguarding, and analyzing data on 47 million vehicles and 12 million crashes, to deliver results about the real-world effectiveness of ADAS. These insights allow partners to make data-driven decisions about enhancements to and investments in advanced driver assistance systems, fostering the safety of US persons traveling by automobile.¹

MITRE has a 50-year history of partnering with federal agencies to apply the best elements of artificial intelligence (AI) and machine learning (ML) while developing and supporting ethical guardrails to protect people and their personal data. Our team's experience with the entirety of the AI/ML adoption and life cycle has strengthened our ability to anticipate and solve future needs that are vital to the safety, well-being, and success of the public and the country. MITRE has deep expertise in systems engineering and integration, having developed architectures for numerous data sharing and analytics platforms. This includes portals, advanced visualizations, and the necessary security controls to enable shared resources. MITRE has earned the reputation and trust of government, industry, and academia as an honest broker, and we stand ready to serve the interests of the Task Force and the needs of the National Artificial Intelligence Research Resource (NAIRR).

Questions Posed in the RFI

a. Vision for the NAIRR. (Chapter 2 of the report)

MITRE recommends that the approach and substance of current NAIRR vision, goals, composition, etc. be strengthened and clarified. Critically, MITRE recommends NAIRR do that through a collaborative approach to strategic planning with the entities that are most likely to be affected by or involved in NAIRR. Based on prior partnerships, MITRE has found that using approaches that engage potential

¹ Partnership for Analytics Research in Traffic Safety. 2022. NHTSA, <https://www.nhtsa.gov/parts-partnership-for-analytics-research-in-traffic-safety>. Last accessed June 29, 2022.

partners and end users and allow them to co-design goals and other foundational aspects of their collaboration and associated partnership operating model is a leading indicator of the eventual success of such partnerships.² To that end, we recommend NAIRR, together with the relevant stakeholders, use a strategic planning framework consistent with the Government Performance and Results Act to strengthen and clarify the NAIRR concept (see Figure 1).

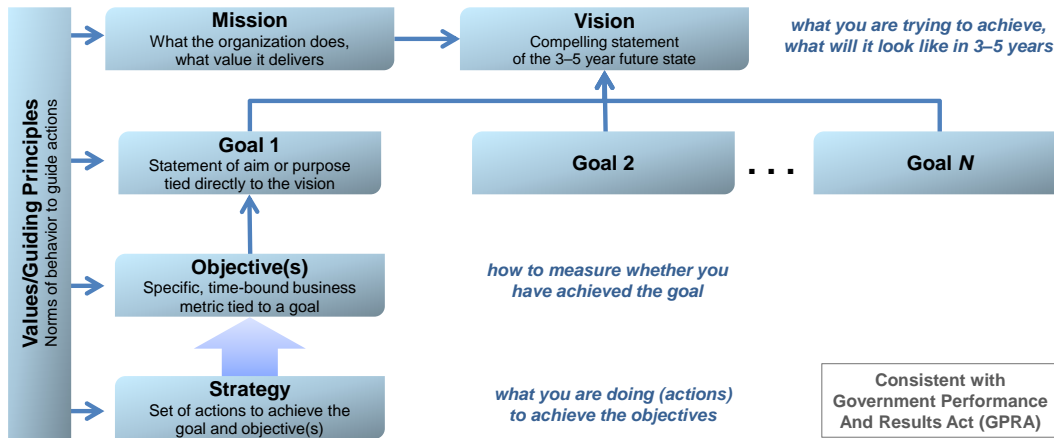


Figure 1 – Strategic Planning Framework with Values/Guiding Principles

Such a co-developed structured planning framework provides:

- A universal and compelling vision for the future of AI research
- A series of goals that collectively enables the vision to be met
- Subordinate objectives and strategies that are specific and time-bound, which both help to drive activities so that they successfully meet the goals and provide NAIRR and the Executive Office of the President (EOP) the ability to measure progress
- A set of values and principles that guides all subsequent activities

We believe that this approach, as opposed to the original construct found within the draft, will be much more comprehensive and greatly enhance NAIRR’s ability to successfully drive intended outcomes.

Specific comments on interim report recommendations:

- [Recommendation 2-2]: Spurring innovation should explicitly name not just “foundational and use-inspired AI research” in general but national objectives-driven AI research, such as U.S. social well-being and equity, health, national security, and the robustness of civic institutions. NAIRR should be strategically filling in the gaps for national objectives that aren’t sufficiently driven by market incentives, or that require cooperation.³
- [Recommendation 2-2]: MITRE concurs that NAIRR should actively seek to increase the diversity of AI researchers since that could lead to new ways of addressing how to reduce the differential performances exhibited by some AI systems. A NAIRR operator with maturity in its own JEDI (Justice, Equity, Diversity, and Inclusion) efforts will likely have more success in increasing the diversity of AI researchers than one with less mature efforts.

² This and subsequent partnership-related recommendations are based on MITRE’s experience- and evidence-based practices gleaned from collaboratively designing and operating many forms of partnerships, including PPPs. MITRE recently published a toolkit with these insights, a portion of which is available at: MITRE’s Public-Private Partnership Accelerator Toolkit (P3TK). 2022. MITRE, <https://ppptoolkit.mitre.org/>. Last accessed June 24, 2022.

³ MITRE Response to OSTP’s RFI Supporting the National Artificial Intelligence Research and Development Strategic Plan. 2022. MITRE, <https://www.mitre.org/sites/default/files/publications/pr-21-01760-16-mitre-response-ostp-rfi-national-artificial-intelligence-research-and-development-strategic-plan.pdf>.

- [Recommendation 2-3]: To the extent that NAIRR involves contributions (e.g., expertise, compute, data, funding) from non-government partners, those contributing partners will likely expect to have a role in shaping the nature of the collaboration, including the NAIRR strategy, composition, and value proposition. NAIRR should extend an approach noted in the Interim Report (tying accountability of the management entity to the Board of Governors and external advisory bodies) to drive partner engagement in earlier stage design, prototyping, and operations.⁴

b. Establishment and sustainment of the NAIRR. (Chapter 3 of the report)

MITRE generally concurs with the recommendations of Chapter 3. The envisioned collaborative and cross-sector nature of NAIRR is similar to other PPPs that we have helped establish and maintain over several decades. We have noted that successful PPPs have three primary characteristics:

1. PPPs are working arrangements based on a mutual commitment—over and above that implied in any contract—between one or more public sector organizations and any other organization(s) outside the public sector to achieve some mutually beneficial outcome.
2. PPPs are collaboratives in which the goals, structure, governance, roles, and responsibilities are mutually determined, and decision-making is shared.
3. PPPs are distinct from traditional contractual arrangements and are rooted in co-creation, co-design, and co-resource mobilization.^{5,6}

Planning for NAIRR can benefit from the lessons learned and proven practices based on innovation-centric and information-sharing PPPs:⁷

- **Innovation-centric PPPs** focus on applied research and development based on the partners' shared interests in reducing their risk and investment to create new intellectual property (IP) and/or cross the chasm of technology adoption. Examples include tech development consortia, national laboratories, government use of cooperative research & development agreements, and similar approaches to stimulate new solutions or markets.
- **Information-sharing PPPs** focus on collaborative data sharing and integrated analyses to produce insights that are otherwise unavailable elsewhere, so that partners can take action on whole-of-nation challenges (e.g., healthcare delivery, quality, and payment integrity; cybersecurity; transportation safety) and realize benefits to their organization and the public.

MITRE recommends that NAIRR consider the following lessons based on MITRE's Public-Private Partnership Accelerator Toolkit given the importance of cross-sector collaboration in achieving NAIRR outcomes.

[Recommendations 3-8, 3-9]: **Establish and reinforce shared decision-making.** PPPs are fundamentally trust-based journeys. In most PPPs, the public partner(s) (i.e., government agencies) share control of the strategy, operations, and decisions with other members of the partnership. This diffusion of power may sometimes cause tension, but achieving whole-of-nation impact requires trust and some give and take.

⁴ Public-Private Partnership Accelerator Toolkit "Value Proposition." 2022. MITRE, <https://ppptoolkit.mitre.org/value-proposition/>. Last accessed June 24, 2022.

⁵ D. Brinkerhoff and J. Brinkerhoff. Public-private partnerships: perspectives on purposes, publicness, and good governance. 2011. Public Administration and Development, https://www.researchgate.net/publication/227724894_Public-private_partnerships_Perspectives_on_purposes_publicness_and_good_governance. Last accessed June 20, 2022.

⁶ Reports – Office of Global Partnerships. 2022. U.S. Department of State, <https://www.state.gov/reports-office-of-global-partnerships/>. Last accessed June 20, 2022.

⁷ MITRE adopted this taxonomy of PPPs to differentiate these newer types of PPPs from traditional, infrastructure-centric PPPs. Most PPPs are for developing (and operating) major public infrastructure such as toll roads or water treatment plants—and are accomplished through a long-term, performance-based government contract that places the management and major share of risk on the private entity.

When the public partner is willing to collaborate, be flexible, and share decision-making authority, there can be large-scale impacts. A successful government partner is prepared to execute many important and distinct roles—champion, funder, recruiter, co-chair of governance bodies, and more—and, equally important, is willing to step back and follow industry/academic partners to advance the shared mission and honor the PPP agreements.

[Recommendations 3-6, 3-20]: **Maximize flexibility.** Enable the PPP to evolve organically with operational flexibility. Successful PPPs allow for the adaptation that happens when you allow smart people collaborating under the right partnership model to respond to emergent challenges, innovations, and their own learning. Agreements and governance that explicitly allow for flexibility and responsiveness to partner input and group-based decisions will serve to advance the shared mission and tap into partner strengths as they collaborate, learn, and adapt together.

[Recommendations 3-2, 3-5, 3-13, 3-14, 4-8]: **Explore and define mutual benefit.** Enable PPP partners—entities that represent groups affecting and affected by the PPP’s work—to explore through prototypes, proofs of concept, and similar lower-risk trials how the PPP will provide each organization benefits that outweigh the cost and risk of their participation. As part of early shaping of collaboratives, MITRE has found it essential that partners gain (at a high level) a clear understanding of the solution the PPP is intended to deliver and a viable idea for how they collaboratively build that solution. Successful PPPs test the following value propositions early in their collaboration: articulate a common understanding of the group’s mission and objectives, aid in recognizing both similar and differentiated benefits of participation, facilitate the buy-in of key partners that will need to help stand up and develop the partnership, and enable the group to share relatable messaging about the PPP’s work when ready to recruit new members. Early conversations among partners will include many thoughts about what the group can accomplish together and what those accomplishments mean for them, the entities they represent, and the overall system of which they are a part. These early thoughts will be tested throughout the proof of concept; only some will emerge as the proven value proposition(s) of the PPP. For an example, see the value proposition exploration process for the MDIAS initiative.⁸

[Recommendations 3-7, 3-8, 3-9, 3-12, 3-15, 3-20]: **Manage expectations.** Cross-sector collaboratives succeed when they openly address the unique needs, interests, and concerns of affected groups. Much like partners co-designing the PPP’s value proposition, MITRE has found that providing stakeholders a safe space to air and collaboratively address needs, risks, and concerns is a leading indicator of PPP success. PPP experts can facilitate business and legal representatives from partner organizations working together to develop mutually satisfactory agreements (and mitigations to address any concerns) about:

- Purpose, governance, roles and responsibilities, and operations of the PPP
- How information is shared, by whom, and when
- Data privacy, security, and permitted uses
- Invention, ownership, and use of intellectual property
- Responses to legal demands for disclosure, Freedom of Information Act
- Measuring outputs and outcomes, learning, and adaptations
- Conflicts of interest, antitrust, unfair competitive advantage, safe harbor, and any other topics specific to the partnership

[Recommendations 3-8, 3-20]: **Collaboratively define guiding principles.** When partners agree to a set of guiding principles—particularly when they co-define those principles—those norms help the PPP

⁸ MDIAS Proof of Concept Overview. 2022. MITRE, https://mdias.org/wp-content/uploads/2022/02/MDIAS-Proof-of-Concept-Fact-Sheet_021122-prs.pdf.

navigate unanticipated issues and gray areas as it operates and adapts to achieve its mission.⁹ Example principles include:

- **Strictly for PPP mission** – Partners share funding, expertise, information, and other in-kind contributions solely for the purposes of the stated mission. Partners will not use this information for unfair competitive advantage, punitive reasons, or any other purpose.
- **Collaborative governance** – Partners co-design the partnership concept, operations, and legal agreements for mutual benefit and to protect partners’ interests and data. Every partner has an equal voice in the consensus-based decision-making of the partnership. Partners commit to working together in good faith to achieve the goals of the partnership.
- **Protection of partner data** – Partners codify legal, security, privacy, ethical, and other expectations to safeguard their information from inappropriate use or disclosure and obtain commitment that all parties will comply. Partners retain ownership and control of their data; if a partner chooses to end participation, that partner’s data is destroyed.
- **Voluntary participation for mutual and public benefit** – Partners voluntarily participate in the partnership predicated on their receiving value from the partnership.
- **Meaningful contributions to transparent operations** – Partners contribute (e.g., time and expertise, information, technology, funding) to the partnership in an equitable and substantive manner, which may vary by task. The PPP operates transparently, with all partners shaping and having access to documented processes, communications, and ways of working together (e.g., collaboration tools and shared AI data/infrastructure).

[Recommendation 3-6, 3-13, 3-14, 4-6]: **Create conditions for trusted collaboration.** If suitably designed, the NAIRR’s management entity could effectively serve as a Trusted Third Party (TTP) among government and non-government partners. Regardless, TTPs can convene partners, nurture relationships, guide collaboration, serve as a trustworthy steward and capable analyst of partner data, and as needed mitigate partner concerns about the inappropriate use of their contributions or unintended effects of their participation in the PPP. This is especially true when the TTP is an independent and objective entity that lacks commercial interests or other potential conflicts of interest. Partners are likely to require an independent and experienced TTP when they seek to mitigate concerns such as:

- **Competitive advantage** – If whatever partners share in good faith is used against them by competitors (in or external to the PPP), or PPP participation affects their market position.
- **Adverse action** – If the agency that regulates them is also a PPP partner, industry partners may be concerned that their participation increases their exposure or could be used for punitive purposes.
- **Commercial conflicts of interest** – If partner data were, for example, to be resold or monetized, or the entities who have access to the data also have a financial interest in the same markets.
- **Protecting IP** – If IP owned by partners is misused or if partners contribute to an invention while participating in the PPP.
- **Exposure** – If the sharing and analysis of data or related PPP activities increases perceived or actual privacy, compliance, or legal risk.¹⁰

⁹ For example, the Department of Transportation’s Partnership for Analytics Research in Traffic Safety (<https://www.nhtsa.gov/parts-partnership-for-analytics-research-in-traffic-safety>). Last accessed June 24, 2022).

¹⁰ Note that successful PPPs requiring a TTP also ensure that the TTP follows the safeguards codified in PPP agreements, such as: properly handles partner-provided data; ensures any partner’s data is not accessible by any other partner including government partners; uses data only for partner-approved purposes; anonymizes data when needed so that individuals are not identifiable in results and results are not attributable to specific partners; is not subject to Freedom of Information Act; and cooperates with the cognizant partner(s) to resist or limit, to the fullest extent permitted by law, any legal process whatsoever demanding the release of any partner information.

[Recommendation 3-5, 3-6, 3-14]: **Explore co-resource mobilization.** As NAIRR demonstrates value to non-government partners, those partners may choose to (further) invest in NAIRR through in-kind and financial contributions. This model of co-resource mobilization mitigates yearly federal appropriations delays and uncertainties. Moreover, this can create a positive feedback loop where more partners are invested in achieving NAIRR outcomes and driving delivery of results, which fosters continued or additional investment. Recommendation 3-1 addresses funding NAIRR through appropriations to multiple federal agencies, which is critical, but overlooks other approaches to long-term sustainability. Private sector contributions can be substantial and be an organic element of NAIRR resourcing and operation. This funding should be accepted under a framework that enables NAIRR to maintain its independence regardless of the source of its finances. NAIRR should explore with potential partners (e.g., cloud services and technology providers) the conditions favoring a broad range of contributions and the optimal resourcing model given capabilities and constraints.

[Recommendation 3-12, 3-11, 3-13]: **Transparent criteria.** NAIRR should provide guidance on transparent, data-driven approaches and methods for selection of applicants competing for NAIRR resources that will achieve NAIRR's objectives of facilitating research with merit while broadening access and participation to underrepresented and underserved researchers and students. A key objective of NAIRR is to broaden access to the resources necessary to conduct AI research to underrepresented and underserved researchers and students, based on research merit. Yet, a tiered structure for NAIRR access based on the cost of resources requested may not be sufficient to provide some amount of higher-cost resources to underrepresented and underserved researchers and students. MITRE recommends NAIRR provide more detailed guidance on the transparent, data-driven approaches for applying selection criteria to groups and individuals competing for NAIRR resources.

[Recommendation 2-1, 2-4]: **Equitable access.** Achieving equitable access to AI resources requires more than making the resources available to the public via application—NAIRR should include a strategy for outreach to underrepresented researchers and students, as well as a plan to provide application support, since better-resourced groups will have more resources and experience in finding these opportunities and preparing the application.¹¹ MITRE identified information access and application support as a common theme across federal agencies' Equity Action Plans. When awarding access to limited resources, NAIRR should explicitly include equity considerations in the selection criteria to counter biases in merit assessment; MITRE has internally studied guidance on social equity considerations in benefit-cost analysis for government grant selection. NAIRR could even consider the option of conducting broader competitions to solve challenge problems. These could be conducted with initial seeding of data, capabilities, and resources. The goal would be to help inform the most viable applicants to receive additional funding and resources to take technical topics further for research.

c. NAIRR resource elements and capabilities. (Chapter 4 of the report)

In answering this question, MITRE provides insights from select existing PPPs that can be advantageous to NAIRR planning.

[Recommendation 4-9, 4-10, 4-11]: **Data sharing.** MITRE concurs that negotiating a data use agreement (DUA) involving any sensitive or proprietary data can be complicated and lengthy. Each party has legal obligations and equities that should be respected. However, there are ways to streamline and standardize these types of negotiations in a way that maintains the integrity of the data sets and trust in the parties' relationship. MITRE has been able to accelerate the approval for data sharing when a small subset of representative partners collaborates on framing a common agreement and the partnership adopts it as a

¹¹ A Framework for Assessing Equity in Federal Programs and Policies. 2021. MITRE, <https://www.mitre.org/sites/default/files/publications/pr-21-1292-a-framework-for-assessing-equity-in-federal-programs-and-policy.pdf>. Last accessed June 24, 2022.

standard DUA template. Some partnerships use an online portal for DUAs where data requesters and data providers only need to enter information in a few fields. When instances require flexibility, MITRE encourages agility and compromise within the bounds of the partnership agreement and guiding principles.

One example of a MITRE-supported information sharing PPP is the ASIAs program. Launched by MITRE and the Federal Aviation Administration in 2007, ASIAs advances aviation safety by leveraging safety data from across the aviation industry to identify emerging systemic risks and to evaluate the effectiveness of deployed mitigations. ASIAs includes government agencies, aviation stakeholder organizations, aircraft manufacturers, and dozens of airlines and corporate operators. The program obtains and fuses data from these partners and other sources so that safety trends can be identified and addressed before accidents or other serious incidents occur. MITRE safeguards this data, which is de-identified, to foster broad engagement and facilitates the data sharing and analysis aspects of ASIAs.¹²

ASIAs is based on the following guiding principles, which foster trust with participating entities:

- ASIAs information is used solely for the identification, monitoring, and mitigation of systemic safety issues.
- Submitted data is not used punitively.
- ASIAs stakeholders voluntarily submit safety-sensitive data.
- Data are de-identified to preserve anonymity.
- Roles and responsibilities of ASIAs stakeholders are developed collaboratively.
- ASIAs data use is transparent to all stakeholders and supporting organizations.

NAIRR may benefit from developing similar guiding principles for its data repositories in collaboration with the data-providing partners. For example:

- NAIRR data repositories are used solely for the creation, testing, and evaluation of AI-enabled capabilities in a manner agreed to by stakeholders.
- NAIRR stakeholders voluntarily submit AI-related data.
- Data are de-identified to preserve anonymity and protected with appropriate controls for sharing.
- Data are characterized for collection methodology and analyzed for bias and ethics, with this characterization contained in data sheets, model cards, and other governance methods.
- Data use is transparent to all stakeholders and supporting organizations.

NAIRR should additionally incentivize the sharing and collection of public-interest datasets in topics not widely available in the commercial and research AI space, aligned with national objectives such as U.S. social well-being and equity, health, national security, and the robustness of civic institutions.

[Recommendation 4-18, 4-19]: **Purpose-suited testbeds.** MITRE concurs that AI comparison testbeds (real-world test, competition, and living laboratory) that are accessible to partnerships and have a low barrier to entry for smaller research entities are an essential element in advancing AI research. Government-funded AI competitions with shared data sets, evaluation protocols, and use cases have successfully driven research in fields such as Natural Language Processing, Computer Vision, Autonomous Systems, and Decision Support.¹³ While these competitions were often focused on a specific domain, they resulted in community data sets and facilitated the sharing of approaches and lessons

¹² Report to Congress: Report on the Status of Aviation Safety Information Analysis and Sharing (ASIAs) Capability Acceleration. 2020. Federal Aviation Administration, https://www.faa.gov/sites/faa.gov/files/2021-11/FAA_Report_on_Aviation_Safety_Information_Analysis_and_Sharing_ASIAs_03312020.pdf.

¹³ Examples of organizations hosting organized AI competitions include: NIST Text REtrieval Conference (TREC)—see <https://trec.nist.gov/>; IEEE International Conference on Acoustics, Speech, & Signal Processing (ICASSP); numerous DARPA programs; and Kaggle competitions—see <https://www.kaggle.com/competitions>. Last accessed June 24, 2022.

learned. NAIRR has an opportunity to advance a more coherent approach to shared testbeds with more principled test and evaluation, which would enable broader and deeper advances in research.

Testbeds should include documentation covering required/provided data sets and evaluations protocols. Data should be characterized using for example Datasheets for Datasets¹⁴ and be properly curated to guard against perpetuating social biases and inequities, for example, by ignoring considerations of underserved populations and diverse demographics.¹⁵ Each AI technology submitted for testing should include a model card documenting how the technology was trained, on what data, for what purpose, how (with results) the technology has been evaluated to date, etc. In hosting testbeds, NAIRR should promote a principled approach to evaluations that whenever possible maps testing in the lab to real-world requirements. This includes the use of evaluation cards that identify the technology tested and document the protocol(s) and data used along with the results of each test conducted. Testbeds should also chronicle lessons learned over time.

Testbeds should provide a means to protect the AI models through their development and training phases, while safeguarding against data exploitation and poisoning. The testbed should be enabled to safeguard against these and other emerging threats to AI that may corrupt the data and models. The testbed too can provide access to sharable insights on ways to design the data to safeguard against these threats. Threat databases are an active area of work with the MITRE ATT&CK¹⁶ knowledge base, and an area of collaboration specifically for AI-enabled systems within MITRE ATLAS.¹⁷

[Recommendation 4-24, 4-25, 4-26]: **Educational Tools and Services.** The Generation AI¹⁸ program is a collaborative program (between MITRE, academia, and private industry) to develop students across the United States into thought leaders who can leverage the power of artificial intelligence and accessible data. Students and faculty in the arts, humanities, and social sciences are tackling real-world challenges alongside their peers in data and computer science. In addition to data, partners share computational notebooks, lecture notes, and homework assignments with one another in the Nexus via our lesson exchange. The goal is to broaden the application of AI and deepen the science—creating a continuous feedback loop that drives innovation and economic expansion. Developing and delivering the program provided several key insights that are useful for planning NAIRR’s educational services:

- Engaging educators and the future AI workforce requires meeting them “where they are.” Rather than focusing on the standard computer science and data science disciplines, the program was able to reach tens of thousands of students across varied disciplines (e.g., fashion design, business) to demonstrate how AI could apply to their differing areas of interest.
- Educators often have limited time and ability to design and implement major changes to their curricula. The greatest opportunities for wider AI education can be delivered by working within existing educational frameworks and integrating smaller, modular lessons that fit to currently defined educational outcomes and lessons.
- As AI educational modules are developed, it is important to have a keen understanding of where educators and students have a level of comfort with coding and other required capabilities. Modules can be designed for educators and students to Use (minimal coding requirements and

¹⁴ T. Gebru, et al. Datasheets for Datasets. 2021. Communications of the ACM, <https://cacm.acm.org/magazines/2021/12/256932-datasheets-for-datasets/fulltext>. Last accessed June 24, 2022.

¹⁵ MITRE’s work on the Maternal Mortality Interactive Dashboard is an example that shows necessity for this consideration. (See <https://www.mitre.org/publications/project-stories/can-data-modeling-and-analytics-help-reduce-pregnancy-related-deaths>. Last accessed June 24, 2022.)

¹⁶ ATT&CK. 2022. MITRE, <https://attack.mitre.org/>. Last accessed June 24, 2022.

¹⁷ MITRE ATLAS (Adversarial Threat Landscape for Artificial-Intelligence Systems). 2022. MITRE, <https://atlas.mitre.org/>. Last accessed June 24, 2022.

¹⁸ Generation AI Nexus. 2022. MITRE, <https://ainexus.org/home>. Last accessed June 24, 2022.

manipulation), Mod (modifying existing code and text for learning), or Create (developing new code for use in AI development).

d. System security and user access controls. (Chapter 5 of the report)

MITRE generally concurs with the findings and recommendations in Chapter 5. Due to the innovative and likely sensitive nature of the data sets maintained by NAIRR, robust cybersecurity and data protection measures are critical to the success of this effort. The threat landscape is rapidly evolving, and NAIRR’s information security team will need to constantly monitor and update security controls to adapt. MITRE concurs with the recommendation that NAIRR adopt a living security plan that evolves with the threat landscape. This security plan should include the adoption of controls from frameworks that partners agree are appropriate for the sensitivity of the data sets; regularly recurring trainings that offer few exemptions from participation and, if so, only based on “testing out;” a clearly defined incident management plan with roles and responsibilities delineated; and an insider threat program that ensures there is no misappropriation or degradation of NAIRR’s systems and data sets.

e. Privacy, civil rights, and civil liberties requirements. (Chapter 6 of the report)

MITRE generally supports the findings and recommendations in Chapter 6. In particular, the themes of transparency, fairness, diversity, adequate privacy engineering, and trustworthiness are essential to establishing a research environment that integrates privacy, civil rights, and civil liberties (P/CRCL) into the NAIRR landscape. These themes should be woven together to form the floor, not the ceiling, of any adequate compliance regime.

MITRE has experience in partnerships that require the development and engineering of privacy frameworks to improve performance and efficiencies while maintaining the public trust in the data sets at hand. This has occurred in engagements across a broad spectrum of partnerships—both public and private entities, large and small. As NAIRR determines how it will implement the recommendations for protecting P/CRCL, MITRE strongly advises it to first develop a framework of how it intends to use the AI data in a way that aligns with certain principles—such as ethical boundaries. This framework would provide a guide to those internally benefiting from NAIRR and offer an understanding to the general public tracking the progress from the outside. This concept is not a new approach. In fact, the Office of the Director of National Intelligence produced the “Artificial Intelligence Ethics Framework for the Intelligence Community” in June 2020.¹⁹ The Department of Defense’s Defense Innovation Unit released its own “Responsible AI Guidelines” in November 2021.²⁰ Furthermore, any mature program that is required to calculate privacy considerations closely follows the path of widely accepted frameworks. NAIRR’s ability to successfully keep P/CRCL at the forefront will require it to establish an AI P/CRCL framework and weave its principles into all its governance and operational documents.

In addition to the AI P/CRCL Framework, NAIRR must be cautious to balance the need for data privacy and protection with the value of the data sets. Because NAIRR is designed to encourage collaboration among a large pool of data users and an even larger pool of data sets, its success in pushing the innovation envelope requires easy access to data. It is a generally accepted privacy principle that data should be shared only with the lowest number of individuals for the least amount of time using it for the least amount of reasons. However, if NAIRR were to follow that model, then its purpose would not be fulfilled. To mitigate this risk, MITRE strongly encourages NAIRR to:

- Incorporate privacy-by-design principles into the various use cases so that each project can include the necessary P/CRCL protections from the beginning of the AI research life cycle.

¹⁹ Artificial Intelligence Ethics Framework for the Intelligence Community. 2020. Office of the Director of National Intelligence, https://www.dni.gov/files/ODNI/documents/AI_Ethics_Framework_for_the_Intelligence_Community_10.pdf.

²⁰ Responsible AI Guidelines – Operationalizing DoD’s Ethical Principles for AI. 2021. Defense Innovation Unit, <https://www.diu.mil/responsible-ai-guidelines>. Last accessed June 24, 2022.

- Mandate AI P/CRCL trainings on an annual basis, with limited exemptions for opting out based on the ability to “test out,” and where necessary require more specific training for researchers and students handling data with higher sensitivity, such as health data.
- Ensure transparency is prominent with all aspects of NAIRR’s responsibilities. This can include hosting public-facing platforms, such as a website and social media accounts, to discuss NAIRR’s ongoing efforts, or it can include standing up a Citizen Advisory Committee to receive feedback from those not official NAIRR researchers/students.
- Engage a diverse group of stakeholders—diverse in technical abilities, professional experiences, educational institutions represented, and personal attributes.
- Conduct random audits of access controls to data sets and oversight of research projects to ensure proper adherence to the AI P/CRCL Framework.

f. Ideas for developing a roadmap to establish and build out the NAIRR in a phased approach, and appropriate milestones for implementing the NAIRR. Including data sets, use cases, and capabilities that should be prioritized in the early stages of establishment of the resource.

As we emphasized above in sections a, b, and c, a partner-driven approach to shaping and operating NAIRR is critical to achieving the intended whole-of-nation impact. MITRE strongly recommends that NAIRR engage the right set of partners to co-create (and routinely revisit and revise) a prioritized roadmap based on their collective strengths and insights. Stakeholders’ buy-in to any roadmap or plan is largely predicated on their degree of involvement in defining it (i.e., seeing themselves in it as a contributor and beneficiary). MITRE also recommends applying organizational change management practices to ensure that stakeholders are ready and supported in accomplishing this journey together.

g. Other areas relevant to the development of the NAIRR implementation plan.

With diversity and growth as a program goal, MITRE recommends proactive outreach that includes key elements, many of which are captured from the MITRE paper *Designing a New Narrative to Build an AI Ready Workforce*.²¹ The government has an opportunity to lead by example in the deployment of responsible AI, and should:

- Define and publicly share its internal governance mechanisms and publicly set expectations with industry partners for deploying AI responsibly.
- Convey legal and ethical accountabilities to the public in a way that describes the responsibility individual decision-makers assume when using any potential system of consequence supporting national security missions.
- Adjust messaging to reflect the values of industry’s founders and modern employees, including preservation of civil liberties, the value of civil service, and humanitarianism.
- Develop opportunities proactively through individual engagements with established interest groups and leverage classic communication methods to shape messaging.

²¹ R. Hodge, et al. *Designing a New Narrative to Build an AI Ready Workforce*. 2020. MITRE, <https://www.mitre.org/sites/default/files/publications/pr-20-0975-designing-a-new-narrative-to-build-an-ai-ready-workforce.pdf>.

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

U.S. Chamber of Commerce Technology Engagement Center

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



June 30, 2022

The Office of Science and Technology Policy &
The National Science Foundation
Attn: Jeri Hessman, NCO
2415 Eisenhower Avenue
Alexandria, VA 22314

Re: RFI Response: National AI Research Resource Interim Report

To Whom It May Concern:

The U.S. Chamber of Commerce's Technology Engagement Center ("C_TEC") appreciates the opportunity to submit comments to the Office of Science and Technology Policy (OSTP) and the National Science Foundation (NSF) Request for Information (RFI) on the "National AI Research Resource Interim Report."¹ C_TEC supports OSTP's and NSF's work to develop a National Artificial Intelligence Research Resource (NAIRR), which "will provide artificial intelligence (AI) researchers and students with access to computational resources, high-quality data, training tools, and user support."²

We wish to provide the below feedback, which further outlines our support on matters we addressed in our previous comments on OSTP and NSF's request for information on the "implementation plan for a National Artificial Intelligence Research resource."³

Federated, hybrid cloud-enabled computing resource:

Our previous comments to the "task force" highlighted the need for the National AI Research Resource to "prioritize developing a hybrid cloud platform that can provide a seamless user experience across multiple clouds."⁴ We strongly believe that only an accessible and easy-to-use hybrid- and multi-cloud computing resource, built on an open architecture that unites both public and private clouds with on-premise resources, can provide the necessary flexibility to provide the scientific community with the resources that are necessary to research at scale.

¹ <https://www.federalregister.gov/documents/2022/05/25/2022-11223/request-for-information-rfi-on-implementing-initial-findings-and-recommendations-of-the-national>

² <https://www.federalregister.gov/documents/2022/05/25/2022-11223/request-for-information-rfi-on-implementing-initial-findings-and-recommendations-of-the-national>

³ https://americaninnovators.com/advocacy/c_tec-comments-on-the-national-artificial-intelligence-research-resource/

⁴ https://americaninnovators.com/advocacy/c_tec-comments-on-the-national-artificial-intelligence-research-resource/

Data and models:

We strongly believe that "Open and Accessible Government Data"⁵ will significantly assist and spur further innovation and breakthroughs within the scientific community, which is why we would like to continue to emphasize the need for the research resource to include high-quality and trusted data sets that the scientific community can utilize. Furthermore, we would like to highlight the importance of the resource to develop pre-trained AI models that researchers can operate in a wide range of disciplines within the AI science technology landscape.

Software and tools:

C_TEC would like to highlight 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 to allow scientists to utilize the research resource efficiently. That being said, we would continue to have concerns regarding the 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.

Education:

We would like to continue to advocate for the need to develop educational materials for researchers who may face challenges in learning how to use the resources. We believe researchers could be assisted by developing training material and workshops to assist them in utilizing the resource and could help accelerate their research. This is why we will continue to advocate for an open line of communication between government, industry, academia, and all stakeholders to learn best practices and necessary training to reduce the skills gap that may be required.

Conclusion

We appreciate NSF's and OSTP's ongoing work to develop NAIRR and efforts to listen to all stakeholders. We highly encourage continuing this critical dialogue with stakeholders to ensure that research resources can be utilized in a way that helps the United States be at the forefront of future scientific discovery. We thank you for considering these comments and would be happy to discuss any of these issues further.

Sincerely,

Director
Chamber Technology Engagement Center

⁵ <https://www.uschamber.com/technology/us-chamber-releases-artificial-intelligence-principles>

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

University of Arizona, CODATA Center of Excellence in Data for Society

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.



June 30, 2022

National Artificial Intelligence Research Resource (NAIRR) Task Force
National Coordination Office for Networking and Information Technology Research and
Development
2415 Eisenhower Avenue
Alexandria, VA 22314 USA

Re: RFI Response: National AI Research Resource Interim Report

Dear Director Parker:

The CODATA Center of Excellence in Data for Society at the University of Arizona (CODATA at UA) is the US-based policy research institute of the International Science Council's Committee on Research Data, or CODATA. We study policies and practices to deliver evidence-based tools and guidance on the use of data assets in society, academia, industry, and government. In particular, we promote the adoption for FAIR data sharing practices that implement Findable, Accessible, Interoperable, and Re-usable data stewardship. FAIR data supports the first principles of innovation—access to tools and knowledge drives inspiration and invention, and improves existing ideas and technologies through ideation and refinement. Knowledge transmission and data sharing are also first principles of basic research, scientific discovery, reproducibility and reproduction. These are the foundational pillars of America's most profound contribution to civilization—a century of scientific discoveries and technologies that have modernized the entire globe and raised billions out of isolation and poverty.

c. NAIRR resource elements and capabilities. Including data, government datasets, compute resources, testbeds, user interface, and educational tools and services. (Chapter 4 of the report)

It is within the context of these first principles and foundational drivers of innovation that we encourage the NAIRR Task Force to adopt the Beijing Declaration on Research Data. This is an existing data sharing framework, designed and refined by hundreds of experts, which takes into account the critical need to protect sensitive data while sharing as much data as

We respectfully acknowledge the land and territories of Indigenous peoples. CODATA at UA is co-located in Tucson AZ on O'odham and Yaqui lands, and in Washington DC on Piscataway, Pamunkey, and Nacotchtank (Anacostia) lands. Work is performed on Yesan (Tutelo) lands on the banks of the New River. Today, Arizona is home to 22 federally recognized tribes, with Tucson being home to the O'odham and the Yaqui. Committed to diversity and inclusion, the University strives to build sustainable relationships with sovereign Native Nations and Indigenous communities through education offerings, partnerships, and community service.



CODATA Center of Excellence in Data for Society

possible toward the benefit of all stakeholders, foreign and domestic. This balance is struck within the Declaration as guidance to make all research data “as open as possible and only as closed as necessary.”

The NAIRRTF may benefit from the research, analysis, and intellectual rigor that the data science community of experts has contributed to the creation of this guidance, which was signed by over a hundred officials and practitioners from government, industry, NGO, and academic sectors. Further, the adoption of existing international frameworks promotes the culture of open science in general, and advances **Recommendation 4-22**, in particular. The adoption of the Beijing Declaration on Research Data supports **Resource Allocation and Sustainment Recommendations 3-10, 3-11, 3-13, and 3-14** by providing a flexible structure for different user, sensitivity, and incentive contexts.

Thank you,

Mercury Fox
Executive Director
CODATA at UA

Paul Uhlir
Consultant
CODATA

Enclosures:

Appendix I: The Beijing Declaration on Research Data

Appendix II: References



APPENDIX I: The Beijing Declaration on Research Data



The Beijing Declaration on Research Data

Preamble

Grand challenges related to the environment, human health, and sustainability confront science and society. Understanding and mitigating these challenges in a rapidly changing environment require data¹ to be FAIR (Findable, Accessible, Interoperable, and Reusable) and as open as possible on a global basis. Scientific discovery must not be impeded unnecessarily by fragmented and closed systems, and the stewardship of research data should avoid defaulting to the traditional, proprietary approach of scholarly publishing. Therefore, the adoption of new policies and principles, coordinated and implemented globally, is necessary for research data and the associated infrastructures, tools, services, and practices. The time to act on the basis of solid policies for research data is now.

The Beijing Declaration is intended as a timely statement of core principles to encourage global cooperation, especially for public research data. It builds on and acknowledges the many national and international efforts that have been undertaken in the policy and technical spheres on a worldwide basis.² These major contributions are listed in the Appendix.

Several emergent global trends justify and precipitate this declaration of principles:

- Massive global challenges require multilateral and cross-disciplinary cooperation and the broad reuse of data to improve coherence concerning recent UN landmark agreements, such as the Paris Climate Agreement, the Sendai Framework for Disaster Risk Reduction, the Sustainable Development Goals (SDGs), the Convention on Biological Diversity, the Plant Treaty, the World Humanitarian Summit, and others. The comprehensive agendas for action provided by these agreements requires access to and reuse of all kinds of data.
- Research and problem-solving, especially addressing the SDG challenges, are increasingly complex and driven by 'big data', resulting in the need to combine and reuse very diverse data resources across multiple fields. This poses an enormous challenge in the interoperability of data and responsible stewardship, with full respect for privacy.
- Rapid advances in the technologies that generate and analyze data pose major challenges concerning data volume, harmonization, management, sharing, and reuse. At the same time, emerging technologies (including machine learning) offer new opportunities that require access to reusable data available in distributed, yet interoperable, international data resources.
- Changing norms and ethics encourage high-quality research through greater transparency, promote the reuse of data, and improve trustworthiness through the production of verifiable and reproducible research results. Increasing the openness of research data is efficient, improving the public return on investment, and generating positive externalities.
- Open Science initiatives are emerging globally, including in less economically developed countries. There consequently are opportunities for these countries to take advantage of technological developments to develop a greater share in scientific production. Without determined action, there is also a risk that the divide in scientific production will widen.

In September 2019, CODATA and its Data Policy Committee convened in Beijing to discuss current data policy issues and developed a set of data policies adapted to the new Open Science paradigm. The Declaration proposed below is the result of that meeting and is now put forward for public review.

¹ In this document we deliberately use the word data very broadly, to comprise data (*stricto sensu*) and the ecosystem of digital things that relate to data, including metadata, software and algorithms, as well as physical samples and analogue artefacts (and the digital representations and metadata relating to these things).

² Europe has been an early mover with its ambitious plans for a European Open Science Cloud, which also has prompted discussions and declarations in many other countries. A collection of major statements of policy principles is listed in the Appendix.



APPENDIX I: The Beijing Declaration on Research Data



The Beijing Declaration on Research Data

The Beijing Declaration supports international efforts to make research data as open as possible and only as closed as necessary. It seeks to make data and metadata Findable, Accessible, Interoperable, and Reusable (FAIR)¹⁸ on a global basis and, wherever possible, automatically processable by machines. Although this Declaration is relevant mostly for research data that are generated through public funding, there are also instances in which privately funded data are made broadly available, in which case these principles would also apply. In addition, data not initially generated for research may be used in research at a later stage. The Beijing Declaration endorses many existing research data policies and management practices that have been promoted by previous declarations and statements, and they are included as references in the Appendix. The participants in the September 2019 policy meeting have produced the following set of ten principles:

1. **Research is increasingly driven by data** that are beyond human processing alone. Researchers therefore should have access to diverse, trustworthy, and reusable sources of data that are readily available and machine actionable. Data stewardship capacity building and comprehensive policies that enable the creation, dissemination, preservation, and above all the **global reuse of data and information** are essential, including sustained support for the required infrastructure and expertise.
2. **Research data have global public good characteristics**. A pure public good cannot be depleted by use (also called non-rivalrous) and cannot be excluded from use. Research data cannot be depleted, but can be restricted in use, although exclusion of reuse by others can be very inefficient and controversial, especially if the data are generated by public funding. **The value of research data increases with use**.
3. Publicly funded research data should be **findable** online to build an **international data commons**. Findable data require comprehensive metadata descriptions and persistent identifier tags, because data that cannot be easily located by potential users—whether by humans or machines—are of limited value. Together, principles three to seven result in “**FAIR**” data (data that are Findable, Accessible, Interoperable, and Reusable)—both for machines and humans.
4. Publicly funded research data are, by default, **in the public interest and should be accessible to the greatest extent possible for international reuse**. They were created or collected on behalf of the public that paid for them, and thus should be as **open as possible and only as closed as necessary**. This is even more important in cases where the data relate to issues covered by the UN landmark agreements.
5. Publicly funded research data should be **interoperable, and preferably without further manipulation or conversion**, to facilitate their broad reuse in scientific research.¹⁹ Software, instruments, and data formats should be well-documented and should not impose any proprietary lock-in that restricts reuse. Data should be described with rich metadata and should use community-recognized terminologies to maximize interoperability and reuse.
6. Despite strong reasons for making research data as open as possible, there are **legitimate reasons to restrict access to and reuse of data**, including interests of national security, law enforcement, privacy, confidentiality, intellectual property, and indigenous data governance, among others. Restrictions should have an express justification and research data **otherwise should be open by default on a global basis**. If the data need to be closed, an effort should be made to provide responsible and proportionately controlled access.
7. **National legislation** that exempts research data from copyright or other intellectual property (IP) protections is one way to enable and support reuse of public data. Another way is for researchers to choose a minimally restrictive and **voluntary common-use license**.
8. Funders of academic and applied research should require the **submission of adequate data stewardship plans**, including clear guidelines for the provision of long-term availability, accessibility, and conditions for reuse. Open data policies should be accompanied by commensurate penalties for noncompliance as well as appropriate incentives.
9. **Activities that address the “divide in scientific production”** between less economically advanced regions and those economies with advanced research infrastructures should include access to publicly funded research data and related information. The wider deployment and access to advanced technical research infrastructures is a necessary, but not sufficient, condition to reduce the divide.
10. **Research data policies should promote the principles in this Declaration and be coordinated internationally**. They should be implemented with clear policy wording and guidelines, specific funding, and a commitment to monitor their impact with the overall objective of building a global FAIR data commons.



APPENDIX I: The Beijing Declaration on Research Data



ⁱⁱⁱ Wilkinson et al 2015, DOI: 10.1038/sdata.2016.18

^{iv} Interoperability of data has technical, semantic, and legal components—all of which need to be addressed successfully to make the data fully reusable.

^v Common-use licenses (such as a Creative Commons CC-BY license or CC0 public domain waiver) preserve some ownership rights while providing access to and reusability of the data. Giving appropriate credit to data providers is essential for promoting data sharing.



APPENDIX I: The Beijing Declaration on Research Data



Appendix

A Selection of Previous Statements and Declarations of Principles on Research Data Policy

1. Australian Code for the Responsible Conduct of Research <https://www.nhmrc.gov.au/about-us/publications/australian-code-responsible-conduct-research-2018> and a set of supporting guidelines including one for Management of Data and Information in Research <https://www.nhmrc.gov.au/file/14359/download?token=0FwepbdZ>
2. Berlin Declaration on Open Access to Knowledge in the Sciences and Humanities. The Max Planck Society. 22 October 2003. Available at: https://openaccess.mpg.de/67605/berlin_declaration_engl.pdf/
3. Budapest Open Access Initiative. Open Society Foundations. 14 February 2002. Available at: <https://www.budapestopenaccessinitiative.org/read/>
4. CARE Principles of Indigenous Data Governance <https://www.gida-global.org/care>
5. Committee on Data for Science and Technology (CODATA) Data Sharing Principles in Developing Countries. Data Sharing Principles in Developing Countries. CODATA. 30 July 2015. Available at: <https://zenodo.org/record/22117#XYxILOzY2w>
6. Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Region, "Open data: An engine for innovation, growth and transparent governance. European Commission. 12 December 2011. Available at https://ehron.jrc.ec.europa.eu/sites/ehron/files/documents/public/2011.ec_communication_open_data_-_engine_for_innovation_growth_and_transparent_governance.pdf
7. Concordat on Open Research Data. (2016). Higher Education Funding Council for England, Research Councils UK, Universities UK, The Wellcome Trust. Available at: <https://www.ukri.org/files/legacy/documents/concordatonopenresearchdata-pdf/>
8. Data Sharing Principles in Developing Countries. CODATA. 18 August 2014. Available at: <https://zenodo.org/record/22117#XRT9NehKhZg/>
9. Declarations in Support of OA (Open Access). Available at: <http://tagteam.harvard.edu/hubs/ostp/tag/oa.declarations>
10. Denton Declaration: An Open Data Manifesto. University of North Texas. 22 May 2012. Available at: <https://openaccess.unt.edu/denton-declaration/>
11. EU Regulation on Copernicus Data, 2014. European Commission. 2014. Available at http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2014.122.01.0044.01.ENG.
12. Executive Office of the President, U.S. Open Data Action Plan. The White House. 9 May 2014. Available at https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/us_open_data_action_plan.pdf
13. Executive Order – Making Open and Machine Readable the New Default for Government Information. The White House. 9 May 2013. Available at <https://obamawhitehouse.archives.gov/the-press-office/2013/05/09/executive-order-making-open-and-machine-readable-new-default-government>
14. FAIR Principles. (2016). GO FAIR. Available at: <https://www.go-fair.org/fair-principles/>
15. G-8 Open data charter. UK Government. 18 June 2013. Available at <https://www.gov.uk/government/publications/open-data-charter>
16. G8 Open Data Charter. UK Cabinet Office. 18 June 2013. Available at: <https://www.gov.uk/government/publications/open-data-charter/>
17. Gates Foundation Policy on Open Access Research. BILL & MELINDA GATES FOUNDATION OPEN ACCESS POLICY. Bill and Melinda Gates Foundation. 1 January 2015. Available at <https://www.gatesfoundation.org/How-We-Work/General-Information/Open-Access-Policy>
18. GEO Data Sharing Principles Implementation. The GEOSS Data Sharing Principles. Group of Earth Observations. 2014. Available at http://www.earthobservations.org/geoss_dsp.shtml
19. Global Biodiversity Information Facility (GBIF), Data Sharing Agreement. New approaches to data licensing and endorsement. GBIF. 22 September 2014. Available at <https://www.gbif.org/news/82363/new-approaches-to-data-licensing-and-endorsement>
20. GODAN Statement of Purpose (2013). GODAN. 2013. Available at <http://www.godan.info/about/statement-of-purpose/>
21. Guidelines on Open Access to Scientific Publications and Research Data in Horizon 2020. European Commission. 11 December 2013. Available at <https://oerknowledgecloud.org/content/guidelines-open-access-scientific-publications-and-research-data-horizon-2020-0>



APPENDIX I: The Beijing Declaration on Research Data



22. ICSU Committee on Freedom and Responsibility in the Conduct of Science (CFRS), Advisory Note on "Sharing Scientific Data with a Focus on Developing Countries. International Science Council. 2011. Available at https://council.science/cms/2017/04/ICSU_CFRS_Advisory_Note_Data_Sharing.pdf
23. Lineamientos de Ciencia Abierta. Colciencias Colombia. 2019. https://www.colciencias.gov.co/sala_de_prensa/participa-en-la-consulta-sobre-lineamientos-politica-ciencia-abierta-para-colombia
24. Lineamientos Generales de Ciencia Abierta. Conacyt México. 2017. <http://www.siiicyt.gob.mx/index.php/normatividad/2-conacyt/1-programas-vigentes-normatividad/lineamientos/lineamientos-generales-de-ciencia-abierta>
25. Manifiesto de Acceso Abierto a Datos de la Investigación Brasileña para Ciencia. Instituto Brasileiro de Informação em Ciência y Tecnología (IBICT). 2016. <http://www.ibict.br/sala-de-imprensa/noticias/item/478-ibict-lanca-manifesto-de-acesso-aberto-a-dados-da-pesquisa-brasileira-para-ciencia-cidada>
26. OECD Principles and Guidelines for Access to Research Data from Public Funding. (2007). OECD. Available at: <http://www.oecd.org/sti/inno/38500813.pdf>.
27. OECD PRINCIPLES AND GUIDELINES. OECD Principles and Guidelines for Access to Research Data from Public Funding. The Organisation for Economic Co-operation and Development (OECD). 12 April 2007. Available at <https://doi.org/10.1787/9789264034020-en-fr>
28. OECD RECOMMENDATION OF THE COUNCIL. OECD Recommendation of the Council for Enhanced Access and More Effective Use of Public Sector Information. The Organisation for Economic Co-operation and Development (OECD). 30 April 2008. Available at <http://www.oecd.org/internet/ieconomy/40826024.pdf>
29. OMB Memorandum to Federal Agencies, "Open Data Policy-Management Information as an Asset.The White House. 9 May 2013 Available at <https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/memoranda/2013/m-13-13.pdf>
30. Open Data in a Big Data World. (2015). International Science Council. Available at: <https://council.science/publications/open-data-in-a-big-data-world/>
31. Open Science Policy Platform Recommendations. European Commission. 22 April 2018. Available at: https://ec.europa.eu/research/openscience/pdf/integrated_advice_opspp_recommendations.pdf#view=fit&pagemode=none/
32. Prepublication of Data Sharing (the Toronto Statement, 2009). National Human Genome Research Institute. 2009. Available at <https://www.genome.gov/27533540/2009-news-feature-prepublication-data-sharing-the-toronto-statement>
33. RECODE Policy Recommendations for Open Access to Research Data in Europe Project. RECODE. 2015. Available at <https://trilateralresearch.co.uk/project/recode/>
34. Reglamento de la ley de AA y Datos Abiertos. Regula el Sistema Nacional de Repositorios Digitales y la que establece el Programa de Datos Abiertos en Ciencia y Tecnología. Government of Argentina. 2017. <https://www.argentina.gob.ar/normativa/nacional/resoluci%C3%B3n-640-2017-277216>
35. Salvador Declaration on Open Access (The Developing World Perspective, 2005). ICML. 21 September 2005. Available at <http://www.icml9.org/meetings/openaccess/public/documents/declaration.htm>
36. Science Creating Solutions—ISC Science Action Plan: 2019-2020. International Science Council. Draft Document. Available at: <https://vgdh.geographie.de/wp-content/docs/2019/02/Draft-Science-Action-Plan-080219.pdf>
37. Science International Accord on Open Data for a Big Data World. Open access to scientific data and literature and the assessment of research by metrics. International Council of Science. 2014. Available at https://council.science/cms/2017/04/ICSU_Open_Access_Report.pdf
38. Sharing Data from Large-scale Biological Research Projects (Ft. Lauderdale Principles, 2003). The Wellcome Trust. 14 January 2003. Available at <http://www.genome.gov/Pages/Research/WellcomeReport0303.pdf>
39. SPARC Europe's Statement on Open Access, 2011. SPARC Europe. 30 May 2011. Available at https://ec.europa.eu/research/science-society/document_library/pdf_06/stakeholder-meeting-minutes_en.pdf
40. Ten Principles for Opening up Government Data.TEN PRINCIPLES FOR OPENING UP GOVERNMENT INFORMATION. Sunlight Foundation. 11 August 2010. Available at <https://sunlightfoundation.com/policy/documents/ten-open-data-principles/>
41. The Australian Open Access Support Group. Open Access Globally. AOASG. 2012. Available at <https://aoasg.org.au/statements-on-os-in-australia-the-world/>
42. The Berlin Declaration on Open Access to Knowledge in the Sciences and Humanities.Max Planck Gesellschaft. 22 October 2003. Available at <http://openaccess.mpg.de/Berlin-Declaration>



APPENDIX I: The Beijing Declaration on Research Data



43. The Bermuda Principles (1996). HUGO. 25 February 1996. Available at http://www.casimir.org.uk/storyfiles/64.0.summary_of_bermuda_principles.pdf
44. The Bethesda Statement on Open Access Publishing (2003). 20 June 2013. Available at <http://legacy.earlham.edu/~peters/fos/bethesda.htm>
45. The Budapest Open Access Initiative. Open Society Foundations. 14 February 2002. Available at <http://budapestopenaccessinitiative.org/>
46. The Declaration on Open Biodiversity Knowledge Management (Bouchout Declaration, 2014). Bouchout 2014. Available at <http://www.bouchoutdeclaration.org/declaration/>
47. The Ghent Declaration (2011). OpenAIRE. 2011. Available at: <https://www.openaire.eu/seizing-the-opportunity-for-open-access-to-european-research-ghent-declaration-published>
48. The Hague Declaration (2014). THE HAGUE DECLARATION ON KNOWLEDGE DISCOVERY IN THE DIGITAL AGE. Liber Europe. 2014. available at: <http://thehaguedeclaration.com/the-hague-declaration-on-knowledge-discovery-in-the-digital-age>
49. The Measures for the Management and Open Sharing of Scientific Data in Chinese Academy of Sciences. 11 February 2019. Available at http://www.ecas.cas.cn/gzwy/ygzzd/201909/t20190917_4552523.html
50. The Vienna Declaration on the European Open Science Cloud. University of Vienna. 23 November 2018. Available at: <https://eosc-launch.eu/declaration/>
51. The White House Office of Management and Budget (OMB) Memorandum on Management of Federal Information Resources, The White House. November 28, 2000. Available at <https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/OMB/circulars/a130/a130revised.pdf>
52. UNESCO Policy Guidelines for the Development and Promotion of Open Access. Policy guidelines for the development and promotion of open access. UNESDOC. 2012. Available at <https://unesdoc.unesco.org/ark:/48223/pf0000215863>
53. White House Office of Science and Technology Policy's Memo "Increasing Access to the Results of Federally Funded Scientific Research". The White House. 22 February 2013. Available at <https://obamawhitehouse.archives.gov/blog/2013/02/22/expanding-public-access-results-federally-funded-research>
54. WMO Resolution 8.1(2), WMO Policy for International Exchange of Climate Data and Products to Support the Implementation of the Global Framework for Climate Services. The United Nation's World Meteorological Organization has issued three resolutions to promote sharing of data. Resolution 25 - World Meteorological Organization. World Meteorological Organization. 1999. Available at http://www.wmo.int/pages/prog/hwrrp/documents/Resolution_25.pdf
55. World Meteorological Organization (WMO)189 Resolution 40, Policy and Practice for the Exchange of Meteorological and Related Data and Products. Report of the Meeting of CCI Task Team on climate aspects of resolution 40. World Meteorological Organization (WMO). 1994. Available at https://library.wmo.int/index.php?lvl=notice_display&id=10466#XYuTS0YzY2x



APPENDIX II: References

CODATA, Committee on Data of the International Science Council, CODATA International Data Policy Committee, CODATA and CODATA China High-level International Meeting on Open Research Data Policy and Practice, Hodson, Simon, Mons, Barend, Uhler, Paul, & Zhang, Lili. (2019). **The Beijing Declaration on Research Data**. Zenodo. Accessed June 30, 2022 at <https://doi.org/10.5281/zenodo.3552330>.

Wilkinson, M., Dumontier, M., Aalbersberg, I. et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* 3, 160018 (2016). Accessed June 30, 2022 at <https://doi.org/10.1038/sdata.2016.18>.

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

University of Southern California (USC) Information Sciences Institute (ISI)

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

RFI Response: National AI Research Resource Interim Report

June 30, 2022

These comments are submitted on behalf of the University of Southern California's Information Sciences Institute.

Response to Recommendation 3-1: Multiple Federal agencies should be funded to cooperatively support NAIRR resources and management, thereby serving the broadest range of research communities and national interests.

There is merit to funding several federal agencies to contribute to NAIRR. This would ensure that all are vested in the effective growth and use of the resource. Indeed, many agencies have served the research community over the years in different ways: the NSF successfully established national supercomputing centers and has supported them for many years to serve the large scientific community; the DoE has provided supercomputing resources for use by the community; and NASA and USGS have been community providers of data. However, a loose independent collection of agencies contributing resources will be a risky approach. Recognizing that selecting a single agency to create the NAIRR would result in a single point of failure for NAIRR, a fully distributed model as stated in the recommendations would be unworkable.

To mitigate this risk, a sensible approach would be to select a small number of agencies to have oversight over the resource. These agencies would have a proven record for serving the community with data and computing resources. An example of a balanced approach could have NSF oversee computing resources for academia, DoE oversee computing resources for government, and NASA oversee data sharing aspects of the NAIRR.

Response to Recommendation 3-4: The day-to-day operations of the NAIRR should be managed by an independent, non-governmental entity with dedicated, expert staff.

Putting NAIRR operations in the hands of a single organization would result in a single point of failure for the entire enterprise. Given that NAIRR is creating a fundamentally new resource like no other, reducing risk should be a major consideration.

To reduce this risk, setting up a consortium of partner organizations would be a reasonable approach. The Open Science Grid (OSG) is an example of a successful consortium to enable sharing of computing resources to serve the community.

Alternatively, a few (3-5) organizations could be responsible for day-to-day operations. Each could be better positioned to serve some segment of the community.

The managing organization(s) of NAIRR should not be in a position of serving many masters. If NAIRR were to be funded by many government agencies then that would make their success more challenging.

Response to Recommendation 3-10: Access to NAIRR resources should be contingent on research project proposal review, be governed by clear use policies and user agreements, and be in compliance with relevant requirements for open sharing of research outputs.

There should be provisions to tie some of the NAIRR allocations more explicitly to research awards. Since research proposals require a description of the resources available to do the work, it would be difficult to obtain funding for ambitious research that ultimately would require resources on the NAIRR. Perhaps a process of pre-approval for using NAIRR for a research proposal if funded would be warranted.

Response to Recommendation 4-1: The NAIRR should coordinate a network of trusted data and compute providers and hosts for a robust, transparent, and responsible data ecosystem.

NAIRR planning should include funding to support AI research on data sharing and data integration. Data sharing has been traditionally challenging, and AI research offers many possibilities to address those challenges by automating data modeling, data integration, and data analysis. In other words, AI presents an opportunity to address traditional challenges in data sharing that have been found in similar efforts at NIH and other agencies. Explicit allocations of funding to support AI research for data sharing should be part of the planning for NAIRR. These funds would not necessarily need to be managed by NAIRR, and could be allocated to funding agencies such as NSF that traditionally fund basic AI research.

NAIRR planning should also include provisions for building on successful AI approaches for data sharing and data integration. WikiData is one such approach that should be considered.

Response to Recommendation 4-5: The NAIRR ecosystem should make the most of community

access by incentivizing the contribution of high-quality data for AI R&D to the federated system.

The Obama administration’s directive for open government data with no additional resources to federal and local agencies resulted in repositories with thousands of datasets that were not very usable. NIH and NSF efforts to create community repositories of shared data have also shown that the effort to share high-quality well-documented data is not always affordable for data providers. It is also not clear that data providers have the knowledge or skills to do the necessary work properly. Therefore, relying on “incentives” for doing what is an enormous amount of work that requires special skills is repeating the same mistakes from the past. It is time to create better approaches for sharing and integrating data through new AI research, so that data sharing becomes cheaper and more scalable. When human effort is needed, there should be a plan to fund the work that includes consideration of the costs and benefits of the target datasets.

The University of Southern California’s Information Sciences Institute (ISI) carries out basic and applied research in artificial intelligence, networks and cybersecurity, high- performance computing, microelectronics, and quantum information systems. Its \$100M annual external funding comes from the NSF, DoD, IC, NIH, DoE, industry, foundations, and other sponsors. ISI is home to the first quantum computer in academia. Part of the USC Viterbi School of Engineering, ISI has more than 400 personnel that includes 28 faculty that advise 65 PhD students. ISI’s Artificial Intelligence Division is one of the largest AI research groups in the U.S. ISI’s AI systems for machine translation, online misinformation detection, and data-centric AI are first-rate and have been deployed to support many parts of DoD and hundreds of law enforcement agencies. Some of ISI's commercial spinoffs were acquired for tens of millions and contribute to a vibrant innovation ecosystem in Southern California.

Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Joseph Wehbe

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

June 29, 2022

Dr. Lynne Parker

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

Re: Request for Information (RFI) on Implementing Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force (Document Number: 2022-11223)

Submitted by:

Joseph Wehbe



“World Economic Forum

Recognized Artificial Intelligence
& Entrepreneurship Expert”

AI Ecosystem Builder

Joseph Wehbe is an American artificial intelligence ecosystem builder. Led the #1 winning team of a Massachusetts Institute of Technology (MIT) Challenge (knowledge-economy) in 2020. He received an AI master’s degree recognized by the leading research institute in Canada in which Dr. Geoffrey Hinton (the Godfather of AI) is the Chief Scientific Advisor. Joseph is also an ambassador for Stanford Women in Data Science in Canada.

Dear Dr. Parker,

There’s a demand for a generation of workers skilled in AI, and it’s my mission to build that by focusing on 3 areas:

1. Operationalizing Federal, State, & Local Govt AI strategies.
2. Building a pipeline of talent & projects as a Government to Grassroots AI value network.
3. Redesigning the entry margin into AI & allowing the non-consumers of AI to participate.

I hereby submit my feedback based on 2+ years of being the class president of an Artificial Intelligence Masters program in Canada led by Dr. Geoffrey Hinton as the Chief Scientific Advisor, and as an American participating in building the Canadian AI ecosystem.

Summary of my NAIRR feedback & level setting:

1

We must define the eligibility of students & researchers who are earning an education in AI. There are only 17 AI focused Master's programs in the US. For us to be inclusive of all students we must redesign the entry barrier to AI for participants in academia, industry, research, entrepreneurship, investors, government, & practitioners. In AI, there's so much public opinion & policy, AI students themselves receive very little say about their own discipline, at the same time, we bear the burden to deliver on the potential of the future while trying to navigate through it all. Let's build an environment that gives back to students what belongs to students, and to seed a culture of learning, innovation, and research.

2

While scientific merit is important as mentioned, educational merit is required of the stakeholders accessing the NAIRR and the AI education development to build a pipeline of AI talent. To address DEIA, we must solve the AI education problem. The barrier now is "those with AI knowledge" and "those without it".

3

We must integrate an infrastructure and software layer to operationalize the NAIRR plan. According to the Global AI Index Report 2022, the US ranks 35th globally in "Operating Agreement", and ranks 17th in "Government Strategy"; this is reflected in our nation's AI strategy execution. The NAIRR plan has the ability to evolve into a Government to Grassroots AI value network for the benefit for American Federal, State, and Local Government stakeholders. We must fix this! Our low ranking in these 2 positions are the basis of all my feedback.

USA Global Index Ranking

*source:
<https://www.tortoisemedia.com/intelligence/global-ai/>



Intelligence | Global AI Index



United States Of America

1	Talent
4	Infrastructure
35	Operating Environment
1	Research
1	Development
17	Government Strategy
1	Commercial
1	Total Rank

FEEDBACK 1

From NAIRR Page 2 Line 4

Going from AI to “organize their days, find the best routes to work and school, select the items they buy, and remind them of upcoming appointments”

→ **[Response by JW]** *As a nation we should think about moving from using AI to “organize our days” to work on projects of National interest building American Dynamism in Aerospace, Defense, Education, Housing, Transportation, Public Safety, Supply Chain, Manufacturing and beyond.*

FEEDBACK 2

From NAIRR Page 1-1

The “growing divide” in computational and data resources

→ **[Response by JW]** *The divide is created by those that “have knowledge about AI” and those that don’t. There isn’t a researcher or AI student in the US that has the AI formal education, proprietary data sources, AI use case knowledge AND has a barrier to start their AI journey. The growing divide is ignited by the knowledge gap. Let us build AI education capacity at the K-12 and university level and that will eliminate the growing divide.*

FEEDBACK 3

From NAIRR page 1-2 “new pathways to participation”

→ **[Response by JW]** *We must redesign the entry margin for the underserved communities to participate. We can’t lower the barrier to AI. AI education is difficult. It must be earned from a university to have educational merit. Bootcamps and certificates are not the solution to finding new pathways to participation.*

FEEDBACK 4

From NAIRR page 1-2

“american researchers to access computational and data resources”

→ **[Response by JW]** *The definition of an American researcher must include a researcher that has an AI education, affiliated to a university in the US, part of an AI degree granting program, affiliated to an AI center of excellence, or an AI research lab. Not every American researcher has AI knowledge to execute, it’s not the NAIRR’s role to educate them, it’s the role of the academic institution they belong to.*

FEEDBACK 5

From NAIRR page 1-3

“National AI Initiative Act of 2020”, the 8-point National AI R&D Strategic Plan

→ **[Response by JW]** *Neither mention the educational merit required of the stakeholders accessing the NAIRR or the AI education development to build a pipeline of AI talent. We must bring back educational merit to AI education and subsequently the stakeholders that benefit from NAIRR.*

FEEDBACK 6

From NAIRR page 1-3 “better understanding the national AI R&D workforce needs”

→ **[Response by JW]** *An AI researcher in the workforce belongs to either a well resourced large scale enterprise (i.e. FAANG or similar company), a well resourced AI non-profit lab (i.e. AI Allen Institute), a venture funded startup, an SMB with limited to no AI expertise on the team, an early stage startup that is not funded nor has the scientific/AI educational merit to work on AI research.*

This group of stakeholders do NOT need access to NAIRR. This is an oversimplification of the landscape, but I argue that the focus for AI R&D workforce needs should be on building AI educational merit for stakeholders from all backgrounds that want to participate in AI.

FEEDBACK 7

From NAIRR page 1-4 & 1-5

“...required elements of the NAIRR roadmap and implementation plan”

→ **[Response by JW]** *There is an infrastructure and software layer missing from operationalizing the plan. According to the AI Index Report 2022, the US ranks 35th globally in “Operating Agreement”, and ranks 17th in “Government Strategy” and this is reflected in this report. We must build an infrastructure and software layer to operationalize the plan as a Government to Grassroots AI value network for the benefit of American Federal, State, and Local Government Stakeholders.*

FEEDBACK 8

From NAIRR page 2-1, Recommendation 2-1

“NAIRR should support early experimentation by students learning how to build and apply AI”

→ **[Response by JW]** *We must define the eligibility of students & researchers who are earning an education in AI. There are only 17 AI focused Master’s programs in the US. A computer science degree that covers AI classes is different from a student earning an AI degree. AI bootcamps and certificates don’t give students practitioner level AI skills with educational merit. For us to be inclusive of all students we must redesign the entry barrier to AI for participants in academia, industry, research, entrepreneurship, investors, government, & practitioners.*

FEEDBACK 9

From NAIRR page 2-2

Increase diversity of talent- “by lowering the barriers of participation for all” regardless of “organizational affiliation”

→ **[Response by JW]** *Means we are removing educational merit if we want security, and accountability...We must redesign the entry margin/barrier to AI not lower the barrier. Organizational affiliation in this case should mean that stakeholders belong to an AI lab, and not any American organization.*

FEEDBACK 10

From NAIRR page 2-3

Mentions “the system should take advantage of existing campus” resources...

→ **[Response by JW]** *We don't need to add new resources, but connect existing campuses and launch AI centers of excellence.*

FEEDBACK 11

From NAIRR page 2-3 (recommendation 2-6: support needs students) point 3

Those studying who are “learning about AI, experimenting with the development of AI models and tools”

→ **[Response by JW]** *The AI programs should be explicit, vetted, recognized by the Department of Education, and have a Chief Scientific Advisor. FYI- there are only 17 AI master's programs in the US.*

FEEDBACK 12

From NAIRR page 2-4 (Access to Startups or SMBs) have federal grants, or SBIR, or STTR

→ **[Response by JW]** *Startups are known to offshore work, we should not grant access. The NAIRR can't control a startup's or SMB's offshore / outsourced resources.*

FEEDBACK 13

From NAIRR page 2-4 (access to Private Sector researchers with Federal funding)

→ **[Response by JW]** *Such researchers should be affiliated to an AI center of excellence, or vetted technology hub / program to prevent bad actors. There are 68 such centers in the US. We can build an AI value network, digitally. Unlike an AI ecosystem, the proposed AI value network is a collection of upstream resources, downstream stakeholders, and subsidiary providers/services supporting a shared business model within our ecosystem. Each node adds value to the end goal of that particular AI stakeholder. This AI value network also serves the non-consumers of AI so that they have a pathway to achieve their goals.*

FEEDBACK 14

From NAIRR page 3-1 Sustainability and long term funding or revenue sources.

→ **[Response by JW]** *By establishing the value network in each community and determining their willingness to pay, we can build several revenue streams and business models.*

FEEDBACK 15

From NAIRR page 3-2 Ownership and Administration “other options may exist”

→ **[Response by JW]** *An infrastructure software layer to operationalize the NAIRR across all stakeholders.*

FEEDBACK 16

From NAIRR page 3-3 The day-to-day operations “employ permanent and diverse staff”

→ **[Response by JW]** *What about qualified AI staff, managers of AI? There’s no mention of such in the report. Can the NAIRR employ enough qualified staff with AI masters degrees?*

FEEDBACK 17

From NAIRR page 3-3 NAIRR management Entity “scientific merit” is mentioned

→ **[Response by JW]** *There should be educational merit to the AI stakeholders accessing. Why should there be educational merit to healthcare/doctors but not for AI practitioners?*

FEEDBACK 18

From NAIRR page 3-4 “resource providers” not duplicate resources

→ **[Response by JW]** *All AI programs have platform companies and resource providers seeking their attention, and offer free resources. We must include them into our value network.*

FEEDBACK 19

From NAIRR page 3-4 “addressing DEIA”

→ **[Response by JW]** *Redesign the entry barrier to participate and increasing the number of AI masters programs and K-12 AI education addresses DEIA. We must be inclusive by increasing access to AI education at the University graduate level.*

FEEDBACK 20

From NAIRR page 3-5 “day to day” operations

→ **[Response by JW]** *There are 8 stakeholders in an AI ecosystem, they should all have a path to contribute, not necessarily all be a user.*

FEEDBACK 21

From NAIRR page 3-5 “Governance and performance”

- **[Response by JW]** *The scientific advisors from the AI labs should all have a seat at the table.*
- *For new research proposals, there should be mechanisms for industry / manufacturing / stakeholders in the heartland and emerging frontier hubs to participate*

FEEDBACK 22

From NAIRR page 3-6 recommendation 3-11 “students, startups”

→ **[Response by JW]** *Access should be given to those with educational merit. Connected to AI programs, labs, or other vetted stakeholder groups.*

FEEDBACK 23

From NAIRR page 3-7 recommendation 3-14 “private entities”

- **[Response by JW]** *The private entities should be connected to an AI lab or center of excellence in their local AI value network*
- *They can contribute data from industry but should be connected to AI centers of excellence at their Local or State Government levels.*

FEEDBACK 24

From NAIRR page 3-7 recommendations 3-15 “NAIRR evaluation methods”

- **[Response by JW]** *Each stakeholder has a different goal in AI, and the outcome / impact on each varies, the measurements should reflect such. There is an 8-stakeholder AI ecosystem model that underpins the performance.*

FEEDBACK 25

From NAIRR page 3-8 recommendation 3-16 “qualified external evaluators”

- **[Response by JW]** *SAME AS ABOVE*

FEEDBACK 26

From NAIRR page 3-9 recommendation 3-19 “publicly accessible platform”

- **[Response by JW]** *The definition of the user roles should all be enabled to AI centers of excellence, accredited AI programs, and not open to the world. A vetted AI stakeholder in the US should belong to one of these institutions. This is an oversimplification but I'm available to explain further.*

FEEDBACK 27

From NAIRR page 3-9, recommendation 3-20 “establish mechanisms” for evaluation...

- **[Response by JW]** *Activity based costing and balance score cards should be integrated into the oversight and transparency to inform improvements to the activities.*

FEEDBACK 28

From NAIRR page 4-1 “...user interface portal”

- **[Response by JW]** *There is an infrastructure and software layer missing for the NAIRR to effectively reach the grassroots. Regardless of the user interface portal, how do vetted AI stakeholders interact through the proposed “user interface portal?”*

FEEDBACK 29

From NAIRR page 4-1 “...set of resources for the AI R&D Community”

- **[Response by JW]** *The eligibility and definition of the AI R&D Community must follow an 8-stakeholder model and exclude startups and those not connected to AI centers of excellence. The reason for startup exclusion is mentioned in this document.*

FEEDBACK 30

From NAIIR page 4-1

“...increasing availability of data... AI-ready data, ethical, privacy, security, and usability”

- **[Response by JW]** *There's no mention of proprietary data, how do we manage the intellectual property for the owner, and provide assurance to the owner that data which could belong to a manufacturer that's willing to share based on their set objective (which was their reason to share it to begin with)?*
- *If all researchers are working on open data sets, who's working on proprietary AI projects? AI researchers must understand the context and domain of the problem they are trying to solve. Hence the 8-stakeholder AI model is required.*

FEEDBACK 31

From NAIIR page 4-1, finding 4-1

“...Rigorous AI R&D is often not possible without high-quality, trusted, dense, and transparent data resources.”

- **[Response by JW]** *I argue that rigorous AI R&D is NOT possible without talent having the educational merit, scientific AI advisors, AI labs, and qualified team support. This component is missing from the report.*

FEEDBACK 32

From NAIIR page 4-2, Finding 4-2

“There are substantial data quality challenges within and across most research domains”

- **[Response by JW]** *I believe there are substantial proprietary data availability challenges within and across most research domains. AI researchers don't understand the business use cases / business value of industry, and industry does not understand the importance of the data. For example: an AI researcher seeking to solve a problem in healthcare, finance, or manufacturing in which they don't have domain expertise. We can and must fix this problem.*

FEEDBACK 33

From NAIIR page 4-2, finding 4-3

“...data curation is a substantial challenge for researchers in all domains”

- **[Response by JW]** *Data curation is not the responsibility of NAIIR. We must design a pathway for the private sector to contribute data via their local AI center of excellence.*

FEEDBACK 34

From NAIIR page 4-2, finding 4-4

“There are substantial costs to combining and linking heterogeneous data.”

- **[Response by JW]** *This is not the responsibility of NAIIR nor its expertise. The concern about R&D data relating to privacy concerns can be managed via the AI centers of excellence.*

FEEDBACK 35

From NAIRR page 4-4, "Recommendation 4-5"

"...incentivizing the contribution of high-quality data for AI R&D to the Federated System"

Recommendation 3-13 on page 3-7

"...incentivize contributions to the NAIRR user community or to the public good"

→ **[Response by JW]** *Rewarding contributors of data "in kind" is beyond the scope of NAIRR. There are too many factors, and considerations to assess. Which can be explained to you at your convenience. By creating a pathway for these contributors via their local AI center of excellence, there must be educational merit to any and all activities relating to data and the proposed federated system.*

→ *Any and all contributors should be vetted stakeholders belonging to an AI center of excellence. Otherwise, we can't control the access to sensitive data. AI is about the people, and each stakeholder has an "individual" behind it who must be vetted and with AI educational merit.*

FEEDBACK 36

From NAIRR page 4-4, recommendation 4-8

"...the NAIRR should provide high-value, core data sets to establish a value proposition and jump-start search and discovery"

→ **[Response by JW]** *This statement is not congruent with previous statement on page 4-3 (recommendation 4-1) that the "sheer volume and variety of data of interest will make it impossible for the NAIRR to curate any of all of it"*

FEEDBACK 37

From NAIRR page 4-3, recommendation 4-1

"...data resources could be contributed by researchers, non-profit or commercial organizations, government agencies, state, local, and/or tribal government, academic institutions, and citizen scientists"

→ **[Response by JW]** *If we treat the field of artificial intelligence with the same academic merit as healthcare, then we can identify who is a stakeholder or citizen scientist. Doctors need a medical degree to practice medicine, but they also have physician's assistants, they have nurses and other medical support specialists. Citizen scientists should belong to an academic institution, AI center of excellence, AI lab, or other vetted AI community/ecosystem. The idea is not to raise the barrier for users/stakeholders, rather the goal is to redesign the entry margin so that everything is done with educational merit.*

FEEDBACK 38

From NAIRR, page 4-5,

"Government data sets... key domains in which the Federal Government could help drive AI-based innovation are transportation, healthcare, and natural hazards research, among many others..."

→ **[Response by JW]** *Each of these domains requires contextual understanding of the AI problem to solve with the said data set owned by the particular Federal Government agency. For example, the proposed AI center of excellence in a region can be supply chain, clean energy or healthcare etc focused to allow the connectivity of resources into the NAIRR system*

FEEDBACK 39

From NAIRR page 4-6, recommendation 4-10, “data generated by Federally funded research”

→ **[Response by JW]** *Despite research that's been federally funded, the day to day employees or stakeholders that are involved in a project might not be American citizens in America. Many projects often outsource/offshore their work, and allowing access to such resources might compromise the integrity of the NAIRR. There are many considerations with a startup being given access that I am ready to share at the appropriate time.*

FEEDBACK 40

From NAIRR page 5-2

“Zero trust architecture presumes that no actor, system, network, or service operating outside or within the security perimeter is trusted”

→ **[Response by JW]** *Why do we adhere to a “zero trust architecture” but do not have a “zero trust AI educated stakeholder” policy? AI is about the people, the technology and algorithms have been commoditized, without the AI educated workforce, we can't undertake cutting edge research and solve real-world problems. Access to NAIRR should be inclusive of those with AI degrees, AI formal education, and other reasons previously mentioned in this submission. We don't want to exclude anyone, at the same time, users should have the AI educational merit from vetted institutions. Let us build the next generation of AI talent so that we remain #1 with AI talent globally.*

CONCLUSION BY JOSEPH WEHBE & FURTHER CONTRIBUTION TO NAIRR

I am ready to serve my country in building the American AI ecosystem. I believe we're at an inflection point in history to execute otherwise we'll lose the AI war. The Government has given us all a platform to act now & thereby ignited a passion in me to believe that there's a call to action to build the next generation of AI talent.

Highlights of AI expertise I offer:

- The benefits of AI ecosystems are distributed unevenly across the US & don't exist in the heartland.
- Dismantle institutional, & systematic barriers that limit opportunities for stakeholders in AI & bring educational merit to the AI workforce.
- Redesign entry margin for stakeholders so the US can build a pipeline of AI talent.

I have both the educational & technical expertise to serve my country in any AI project that will keep the US as a world AI leader. “Until the mayor or superintendent in small town New Jersey understands they must introduce an AI K-12 curriculum, we have a lot of work to do.”
-Joseph Wehbe

Your faithfully,

Joseph Wehbe
“AI Ecosystem Builder”

“The recipe is straightforward, let us invest in AI Education, AI Research & Development.”-JW



Request for Information (RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response

Robin Wieder

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views 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 the responses and external links included in this document.

NAIIR Response – Robin Wieder

The government should make direct investment to quantify the multi-dimensional tradeoff between accuracy, fairness, privacy, robustness, explainability, and other societally desirable parameters in machine learning. It should not just study the tradeoffs between any two of these parameters but work to understand the Pareto frontier among all of them simultaneously (or as many as possible) through real-world measurement, simulation, and theoretical study.

Robin Wieder