


*Envisioning a National Artificial Intelligence
Research Resource (NAIRR):
Preliminary Findings and Recommendations*

AN INTERIM REPORT BY THE NAIRR TASK FORCE

May 2022



Dear Mr. President and Honorable Members of Congress:

Artificial Intelligence (AI) is no longer merely the subject of science fiction or experimentation in research laboratories; AI is increasingly integrated into our everyday lives. Many Americans rely on AI to organize their days, find the best routes to work and school, select the items they buy, and remind them of upcoming appointments. AI is driving American ingenuity and helping to develop solutions to the big challenges that our world faces, from optimizing food production to addressing climate change to curing cancer. Once the work of a few, AI is now integral to most areas of research and nearly all economic sectors, and its responsible deployment is crucial to our Nation's long-term economic competitiveness and security.

At the same time, pathways into AI research are too often accessible only to a limited few. Much of today's cutting-edge AI research relies on access to large volumes of data and advanced computational power, which are often unavailable to many researchers beyond those at large technology companies and well-resourced universities. This lack of access can lead to communities—particularly those that have long been underrepresented or underserved—being left out of the AI research and development process. In turn, this lack of diversity in the design of novel AI systems limits the breadth of ideas incorporated into AI innovations and contributes to biases and other systemic inequalities. Simply put, it matters who is developing AI technologies and which principles that developer is embracing to examine the broader impacts.

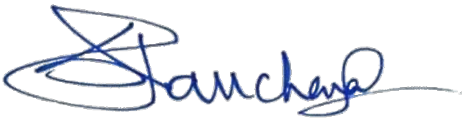
For the United States to sustain its leadership in AI and ensure that AI benefits all Americans, it is essential that the full and diverse talent of the Nation contributes to the AI innovation ecosystem. This requires expanding access to the necessary resources.

Congress recognized this need when, as part of the National AI Initiative Act of 2020, it directed the National Science Foundation (NSF), in consultation with the White House Office of Science and Technology Policy (OSTP), to establish a task force to create a roadmap for a National AI Research Resource (NAIRR)—a shared research infrastructure that would provide AI researchers and students with significantly expanded access to computational resources, high-quality data, educational tools, and user support. This increased access can also help diversify the workforce developing these technologies.

OSTP and NSF formally launched the NAIRR Task Force in June 2021, appointing 12 leading experts equally representing academia, government, and private organizations. Since its establishment, the Task Force has held 7 public meetings, engaged with 39 experts on a wide range of aspects related to the design of the NAIRR, and considered 84 responses from the public to a request for information. We extend our gratitude to the members of the Task Force who have donated an extraordinary number of hours of their time to this effort, as well as to the many members of the public who have contributed their expertise and provided inputs to the Task Force.

This interim report represents the initial results of that work. It sets forth the Task Force's vision for the NAIRR, providing findings and recommendations on a variety of topics as directed by Congress. Moving forward, the Task Force will continue to seek input from the public as it deliberates on an actionable implementation roadmap to be released about 6 months from now. We also welcome your feedback.

Sincerely,

A handwritten signature in blue ink, appearing to read "Sethuraman Panchanathan". The signature is fluid and cursive, with a large initial "S" and a long horizontal stroke at the end.

Sethuraman Panchanathan, Ph.D.
Director,
National Science Foundation

A handwritten signature in blue ink, appearing to read "Alondra Nelson". The signature is cursive and elegant, with a prominent "A" and "N".

Alondra Nelson, Ph.D.
Deputy Assistant to the President
Performing the Duties of the Director,
Office of Science and Technology Policy

Executive Summary

Artificial Intelligence (AI) is transforming our world. The field is an engine of innovation that is already driving scientific discovery and economic growth, and an integral component of solutions that stand to impact everything from routine daily tasks to societal-level challenges. To realize this promise, we must ensure that everyone throughout the Nation has the ability to pursue cutting-edge AI research. Yet progress at the current frontiers of AI is often tied to access to large amounts of computational power and data. Such access today is too often limited to large technology companies and well-resourced universities. This growing resource divide has the potential to adversely skew our AI research ecosystem, and, in the process, threaten our Nation's ability to cultivate an AI research community and workforce that reflect America's rich diversity—and harness AI in a manner that serves all Americans.

Given this current landscape, the National AI Research Resource (NAIRR) Task Force concludes that coordinated action is critical. As a nation, we must come together to expand access to the resources that fuel AI, providing pathways for more Americans to pursue AI research and development (R&D) and access state-of-the-art resources. These pathways would broaden the range of researchers involved in AI, growing and diversifying approaches to and applications of AI and opening up opportunities for progress across all scientific fields and disciplines, including in critical areas such as AI auditing, testing and evaluation, trustworthy AI, bias mitigation, and AI safety. Increased 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.

The vision for a NAIRR laid out in this interim report of the NAIRR Task Force focuses on developing a concept that would meet this national need through a shared research cyberinfrastructure connecting researchers to the resources and tools that fuel AI R&D. The Task Force presents a path for doing so in a manner that builds from existing Federal investments; designs in protections for privacy, civil rights, and civil liberties; and promotes diversity and equitable access. If successful, the 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.

This report provides the Task Force's general vision for a NAIRR along with a preliminary set of findings and recommendations for the design of the NAIRR architecture, resources, capabilities, and uses:

The strategic objective for establishing a NAIRR is to strengthen and democratize the U.S. AI innovation ecosystem in a way that protects privacy, civil rights, and civil liberties. To achieve this objective, the NAIRR should be designed to help achieve four primary goals for AI R&D: (1) spur innovation, (2) increase diversity of talent, (3) improve

capacity, and (4) advance trustworthy AI. The NAIRR should be formulated as a federated cyberinfrastructure ecosystem, accessible through an integrated portal and run by a single management entity, with governance and external advisory bodies to provide oversight and guidance. The NAIRR should support the needs of America's AI researchers and students from diverse backgrounds who are pursuing foundational, use-inspired, and translational AI research.

The NAIRR requires a collaborative effort, with contributions from a collection of Federal agencies that represent AI stakeholders. Today, multiple Federal agencies serve to advance the Nation's AI capability. Therefore, multiple Federal agencies should be funded to work cooperatively to support NAIRR resources and management, with representatives from each of the agencies engaging with the NAIRR management and administration to provide governance and oversight. These Federal agencies should make available via the NAIRR new or existing cyberinfrastructure resources relevant for AI R&D, to include data, compute, and testbeds. Partner resource providers, including private sector providers, should also have the ability to make a broad variety of resources available through the NAIRR user access portal.

The NAIRR requires that key resource elements, including curated data sets, computational power, testbeds, and educational tools and services, be accessible in user-friendly ways. The NAIRR should coordinate a network of trusted data and compute providers and hosts to foster a robust, transparent, and responsible data ecosystem. The NAIRR should make the most of community access by incentivizing the contribution of high-quality data for AI R&D to the federated system, and establishing a value ecosystem around data that can be used for AI and to support data search and discovery. In providing access to government data, the NAIRR management entity should build on and leverage existing Federal data sharing efforts and explore facilitating access to three types of government data: statistical data, administrative data, and data generated by federally funded research. Privacy should be protected by following the "Five Safes" framework for safe use (safe projects, safe people, safe data, safe settings, and safe outputs).

As part of the ecosystem of resources, the NAIRR should provide access to a federated mix of on-premise and commercial computational resources, including conventional servers, computing clusters, high-performance computing (HPC), and cloud computing, and also support access to edge computing resources and testbeds for AI R&D. Software, training, and educational resources should be made available to support a diverse set of users with varying levels of proficiency. The NAIRR should provide multiple levels of user support including help desk and solution consulting, and incentivize community-based support.

The day-to-day operations of the NAIRR require an independent, non-governmental entity with dedicated, expert staff. In addition to managing the cyberinfrastructure, resource allocation, user support, and security, this management entity should be explicitly charged with addressing diversity, equity, inclusion, and accessibility issues related to NAIRR access and NAIRR-supported AI R&D. A formal charter and associated policies should

governance management and administration of the NAIRR, with performance of the NAIRR overseen by a board of governors complemented with mechanisms for external advice, oversight, and evaluation. The NAIRR charter should also establish mechanisms by which evaluation can directly inform adjustments to the strategic goals, operational functions, and resource capabilities of the NAIRR to optimize impact of investments and meet changing needs.

Resource allocation processes for access to NAIRR resources must provide accessible on-ramps for the range of expected users, be lightweight where possible, and be as inclusive as feasible. Costs of allocations should follow a tiered model, where some resources are fee based and some are provided at no cost. The lowest cost resources should be made widely accessible to the broadest range of users. Researchers should in general be able to apply and compete for allocations of limited resources, with decisions informed by research merit and the objective of broadening access and participation. The framework for NAIRR resource allocation should be designed to incentivize contributions to the NAIRR user community or to the public good.

The NAIRR must secure the research, data, resources, and safety of its users while balancing attributes of usability that will be needed to increase participation in AI research. AI research is a high-value target because of its nature as an asset to economic growth and national security, necessitating a focus on providing an effective and trusted research environment.

Fostering an open research environment has tradeoffs with providing secure access to high-value information and resources. Security needs will likely evolve rapidly. If the decision is made for the NAIRR to include sensitive or confidential data among its resources, the NAIRR management entity should embrace a tiered access model to accommodate heterogeneous security needs and employ dedicated technical security staff experts who can keep up with evolving requirements. In addition, NAIRR staff and users should be required to fulfill regular and continuous hands-on security training, and routine monitoring and update of the system and associated security controls should be included as part of NAIRR management activities.

The NAIRR can set the standard for responsible AI research through the design and implementation of its governance processes. To earn and maintain public trust, the NAIRR management entity will need to transparently demonstrate how research using the NAIRR is being reviewed, approved, and performed in a way that meets the expectations of the public.

In support of these principles, the NAIRR should establish an ethics review process to vet all resources included in the system and the research performed within it. NAIRR users should be required to complete regularly updated ethics training modules before being granted access to the NAIRR. In addition, NAIRR resources should be specifically allocated to support research on AI trustworthiness, and to develop best practices for responsibly working with data and models.

Next Steps for the Task Force

The preliminary recommendations presented in this interim report outline the NAIRR Task Force's current vision for the NAIRR and the impact that the NAIRR could have on America's AI research ecosystem. Upon release of this report through June 30, 2022, the NAIRR Task Force will solicit public feedback on the proposed features of the NAIRR through a Request for Information (RFI). To share input, the Task Force invites all interested stakeholders to follow the RFI guidance provided in [Federal Register Notice 2022-11223](#). In addition, a public listening session will be held on June 23, 2022, to provide another opportunity for the public to provide input. Details on how to participate are provided in [Federal Register Notice 2022-11222](#). In the months ahead, the Task Force will integrate this public feedback with other inputs, and further deliberate to arrive at a final report outlining a detailed roadmap and implementation plan for the NAIRR. This final report is anticipated for release by the end of this year.

Contents

1. Introduction	1-1
Unlocking the Potential of AI	1-1
Federal Activities to Advance AI R&D	1-3
Charge to the NAIRR Task Force	1-4
Task Force Approach	1-5
2. A Vision for the NAIRR	2-1
Strategic Objective of the NAIRR	2-1
Goals for Establishment and Sustainment of the NAIRR	2-1
Composition of the NAIRR	2-2
Access to the NAIRR	2-3
NAIRR User Base	2-3
3. Establishing and Sustaining the NAIRR	3-1
Agency Funding, Roles, and Responsibilities	3-1
Ownership and Administration	3-2
Governance and Oversight	3-4
Resource Allocation and Sustainment	3-6
NAIRR Performance Indicators and Metrics	3-7
4. NAIRR Resource Elements and Capabilities	4-1
Data	4-1
<i>Findings</i>	4-1
<i>Recommendations</i>	4-3
Government Data Sets	4-5
<i>Recommendations</i>	4-6
Compute Resources	4-7
<i>Findings</i>	4-7
<i>Recommendations</i>	4-8
Testbeds	4-10
<i>Findings</i>	4-10
<i>Recommendations</i>	4-12
User Interface	4-13
<i>Findings</i>	4-14
<i>Recommendations</i>	4-14
Educational Tools and Services	4-16
<i>Findings</i>	4-16
<i>Recommendations</i>	4-17
5. System Security and User Access Controls	5-1
<i>Findings</i>	5-1
<i>Recommendations</i>	5-2
6. Privacy, Civil Rights, and Civil Liberties	6-1
<i>Findings</i>	6-1
<i>Recommendations</i>	6-2
7. Next Steps for the Task Force	7-1
Appendix A. Definitions	A-1
Appendix B. Briefers to the Task Force	B-1
Appendix C. Public Input Provided in Response to the Federal Request for Information	C-1

Appendix D. Subject Matter Experts Consulted by TF Members..... D-1
Appendix E. NAIRR Task Force Staff and ContributorsE-1
Appendix F. Abbreviations.....F-1
Appendix G. Notes G-1

Task Force Membership

LYNNE E. PARKER (CO-CHAIR)

*Deputy United States Chief Technology Officer and Director of the National AI Initiative Office,
White House Office of Science and Technology Policy*

ERWIN GIANCHANDANI (CO-CHAIR, JULY 2021–OCTOBER 2021)

*Assistant Director for Technology, Innovation and Partnerships,
National Science Foundation*

MANISH PARASHAR (CO-CHAIR, BEGINNING OCTOBER 2021)

*Office Director of the Office of Advanced Cyberinfrastructure,
National Science Foundation*

DANIELA BRAGA

Founder & CEO, Defined.ai

MARK E. DEAN

OREN ETZIONI

CEO, Allen Institute for AI

JULIA LANE

Professor, New York University

FEI-FEI LI

*Sequoia Professor of Computer Science at Stanford University and Denning Co-Director of the
Stanford Institute for Human-Centered AI (HAI)*

ANDREW MOORE

VP & General Manager, Google Cloud AI & Industry Solutions

MICHAEL L. NORMAN

Distinguished Professor, University of California, San Diego

DAN STANZIONE

*Executive Director, Texas Advanced Computing Center/Associate
Vice President for Research, The University of Texas at Austin*

FREDERICK H. STREITZ

*Chief Computational Scientist, Lawrence Livermore National Laboratory,
U.S. Department of Energy*

ELHAM TABASSI

*Chief of Staff, Information Technology Laboratory,
National Institute of Standards and Technology*

1. Introduction

Unlocking the Potential of AI


The term "Artificial Intelligence" (AI) refers to a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments (see also [Box 1](#)).¹ AI systems can be applied to tasks spanning diverse areas, including planning and optimization, perception and vision, modeling and simulation, natural language understanding, robotic process automation, and prediction.

Driven by U.S. Government support and innovation in the public sector, private sector, and academia, the field of AI has made remarkable advancements since it emerged as a discipline in 1956. Initial activity was tempered by a lull in the 1970s, as early progress failed to meet expectations, but was followed by a reinvigoration in the 1980s due to increased commercial interest and the development of the Internet. After a second, so-called "winter," increased computational power and data availability in the 2010s enabled significant advances in the application of deep learning and neural networks to key areas such as speech recognition, visual object recognition, and machine translation.

AI research has reached a new era marked by the wide, practical deployment of various AI technologies across a range of industries and contexts and the existence of large foundational models that can be leveraged and adapted to apply to a variety of tasks.² AI technologies have supported scientific and technological breakthroughs in areas such as protein folding, nuclear fusion, and even coding and theorem proving.³ Many recent advances in AI have relied on large and growing amounts of data and computational power. In some subfields of AI, these resource requirements have grown so great that only the very largest private firms can participate.⁴ In other less resource-intensive subfields, the cost of access to cutting-edge data and computation resources has concentrated research activity among large private-sector firms and the most well-resourced universities.⁵

Currently, access to the computational and data resources that fuel much of today's AI is concentrated in those large private-sector firms, well-resourced universities, and national laboratories, creating a growing divide.⁶ This disparity in availability of AI research resources affects the quality and character of the U.S. AI innovation ecosystem, contributing to a "brain drain" of top AI talent from the vast majority of academic and research institutions to a small set of large corporations.⁷ Such trends have adverse implications for the Nation's capacity to train the breadth of diverse talent required to support future U.S. competitiveness and innovation. The concentration of academic AI research at elite universities also limits the diversity of those studying and designing AI systems. For example, traditionally underserved communities lack sufficient representation and pathways to participation in the field.⁸ Of the new U.S. resident AI PhDs in 2020, approximately 51 percent were non-Hispanic White, 30 percent Asian, 7 percent Hispanic, and 2 percent Black or African American (see [Figure 1](#)). The percentage of AI PhDs awarded to Hispanic and Black or African American students decreased relative to 2010.⁹ Gender diversity in AI has demonstrated little change over the past decade; according to one estimate, only

about 20 percent of both the AI PhD and computer science PhD graduates in North America in 2020 were female.¹⁰ This lack of diversity contributes to instances of development of AI tools and approaches that perpetuate bias, limits the breadth of ideas incorporated into AI innovation, and perpetuates systemic inequalities.¹¹ A related trend is a drop in the number of new AI PhDs to U.S.-based individuals, further impacting representation in the field for underserved communities and illustrating the need for new pathways to participation.¹²



Source: 2022 AI Index Report

Figure 1. Demographics of 2021 AI PhD graduates

AI holds great potential to be an engine of innovation, providing further advances in science, economic growth, national security, and the ability to meet pressing societal challenges.¹³ For example, AI breakthroughs could accelerate progress across a broad spectrum of actions needed to build a more sustainable future—from mitigation of greenhouse gas emissions to the development of data-driven strategies for conservation efforts, to automated solutions for managing consumption, to the invention of new clean energy sources and materials. Realizing such benefits will rely on the ability of American researchers to access computational and data resources, and on building a diverse talent pool to research new algorithms, engineering methodologies, and inspiring use cases.

An expansion of access to the cyberinfrastructure—including computer hardware, algorithms, data, software, services, networks, and expertise—necessary to conduct AI research and development (R&D) would present an enormous opportunity for the United States. It could power efforts to help broaden the range of researchers involved in AI, grow and diversify approaches to and applications of AI, and open up opportunities to advance R&D in AI. Further, it would facilitate use-inspired and translational AI R&D across all scientific fields and disciplines, as well as in critical areas such as AI auditing, testing and evaluation, bias mitigation, and security. Increased access and the resulting diversity of perspectives may also, in turn, lead to new, transformative ideas that would not otherwise emerge, and could advance U.S. economic competitiveness, improve quality of life, and strengthen national security.

Federal Activities to Advance AI R&D

In January 2021, as part of the National Artificial Intelligence Initiative Act of 2020,¹⁴ Congress established the National AI Initiative to ensure that the Nation continues its leadership in AI R&D; leads the world in the development and use of trustworthy AI systems across public and private sectors; prepares the current and future U.S. workforce for integration of AI systems across all sectors of the economy and society; and coordinates AI research, development, and demonstration activities among Federal agencies. One of the activities assigned to the Initiative to achieve these goals is sustained and consistent support for AI R&D through grants, cooperative agreements, testbeds, and access to data and computing resources. This legislation also provided a statutory definition of AI, provided in [Box 1](#), which is used for the purposes of this report. (See also [Appendix A](#), which provides definitions of common terms used in this report.)

Box 1. Definition of Artificial Intelligence, as used in this report (from the National AI Initiative Act of 2020, 15 U.S.C. § 9401(3))

The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments.

Artificial intelligence systems use machine and human-based inputs to—

- (A) perceive real and virtual environments;
- (B) abstract such perceptions into models through analysis in an automated manner; and
- (C) use model inference to formulate options for information or action.

Congress called for these efforts to build upon the *National AI R&D Strategic Plan*,¹⁵ first released in 2016 and updated in 2019, which highlights eight key elements of a national AI R&D strategy: (1) make long-term investments in AI research, (2) develop effective methods for human-AI collaboration, (3) understand and address the ethical, legal, and societal implications of AI, (4) ensure the safety and security of AI systems, (5) develop shared public data sets and environments for AI training and testing, (6) measure and evaluate AI technologies through standards and benchmarks, (7) better understand the national AI R&D workforce needs, and (8) expand public-private partnerships to accelerate advances in AI. A 2020 report of the Select Committee on Artificial Intelligence of the U.S. National Science and Technology Council also outlines recommendations for accelerating the use of cloud computing resources for federally funded AI R&D.¹⁶

In alignment with these strategies, the National AI Initiative Act of 2020 also established the National Artificial Intelligence Research Resource (NAIRR) Task Force (TF) to "explore the feasibility and advisability of developing a NAIRR and develop a roadmap and implementation plan for its establishment" (15 U.S.C. § 9415).¹⁷ A NAIRR—that is, a widely accessible, AI-specific research cyberinfrastructure system, as defined in [Box 2](#)—has the potential to support several elements of the *National AI R&D Strategic Plan* and help to build a stronger, more inclusive U.S. AI R&D ecosystem. Note that, throughout this report, "AI R&D" is inclusive of foundational AI R&D, use-inspired AI R&D, and translational AI R&D. That is, the NAIRR has relevance not only for researchers advancing the field of AI itself (foundational research) but also for those who are advancing AI with a use case in mind (use-inspired research), as well as for those translating AI discoveries and innovations to the market and society (translational research).

Box 2. Definition of NAIRR (15 U.S.C. § 9415(g)(1))

The terms "National Artificial Intelligence Research Resource" and "Resource" mean a system that provides researchers and students across scientific fields and disciplines with access to compute resources, co-located with publicly-available, artificial intelligence-ready government and non-government data sets and a research environment with appropriate educational tools and user support.

Charge to the NAIRR Task Force

Congress charged the TF with proposing a national solution to provide researchers and students across scientific fields and disciplines with access to data and computing resources for AI R&D, along with appropriate educational tools and user support. Specifically, Congress directed the TF to develop a roadmap and implementation plan for establishing the NAIRR.

The TF was launched on June 10, 2021, as a Federal Advisory Committee co-chaired by the National Science Foundation (NSF) and the White House Office of Science and Technology Policy (OSTP), and includes representatives from the U.S. Government, academia, and the private sector. Its members' expertise spans foundational, use-inspired, and trustworthy AI R&D, as well as research cyberinfrastructure. This interim report constitutes the TF's first deliverable, pursuant to its Congressional mandate.

Congress specified that the NAIRR roadmap and implementation plan address nine key dimensions, as stated in [Box 3](#). The TF activities are bounded to developing recommendations and proposing a roadmap and implementation plan for a NAIRR to the President and to Congress. The TF will conclude its work within 90 days after submission of its final report; the TF itself will not execute any of its recommendations, nor will it be involved in the administration of a future NAIRR.

Box 3. Required elements of the NAIRR roadmap and implementation plan (15 U.S.C. § 9415(b))

- (1) IN GENERAL.—The Task Force shall develop a coordinated roadmap and implementation plan for creating and sustaining a National Artificial Intelligence Research Resource.
- (2) CONTENTS.—The roadmap and plan required by paragraph (1) shall include the following:
 - A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource, and metrics for success.
 - B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including—
 - i. an appropriate agency or organization responsible for the implementation, deployment, and administration of the Resource; and
 - ii. a governance structure for the Resource, including oversight and decision-making authorities.
 - C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources.
 - D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to computing resources for researchers across the country, including scalability, secured access control, resident data engineering and curation expertise, provision of curated data sets, compute resources, educational tools and services, and a user interface portal.
 - E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource.

- F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its research and a recommendation for a framework for the management of access controls.
- G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research.
- H. A plan for sustaining the Resource, including through Federal funding and partnerships with the private sector.
- I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities and milestones to implement the Resource.

Task Force Approach

Between its launch date and the time of this report's publication, the TF convened 7 virtual public meetings to discuss and deliberate on key NAIRR uses, potential impacts, system requirements, and design elements. At these meetings, the TF heard from 39 expert briefers and panelists to augment the members' own expertise, and to ensure that a diversity of perspectives and experiences was considered in TF discussions and deliberations. Topics addressed by these invited experts spanned user perspectives; privacy, civil rights, and civil liberties requirements; broadening access and participation; diversity, equity, and inclusion; user resources; data needs and challenges; governance and administration models; related Federal programs; and the value proposition and intended outcomes of a NAIRR. Panelists were drawn from a range of groups and organizations, including national laboratories, nonprofits, civil society organizations, companies, government agencies, colleges, and universities. See [Appendix B](#) for a complete list of invited panelists who provided input to the TF.

The TF also reviewed 84 public responses to a July 2021 request for information (RFI) regarding key aspects of the NAIRR. These responses reflect feedback from individuals (ranging from academics to government researchers to business leaders to interested members of the public), groups, and organizations (spanning nonprofits, civil society groups, research organizations, small and large businesses, and the Federal Government). For a full list of respondents and a link to the RFI responses, see [Appendix C](#).

The TF leveraged its members' expertise to develop proposals for key aspects of the NAIRR. In this context, TF members also engaged with additional outside subject matter experts (see [Appendix D](#) for a complete list of experts consulted) and generated draft findings and recommendations that seeded the development of this interim report.

This document provides the TF's general vision for a NAIRR along with a preliminary set of findings and recommendations for the design of the NAIRR architecture, resources, capabilities, and uses. Concurrently with this report's publication, the TF is issuing a second RFI and holding a listening session to solicit public feedback on the TF's preliminary findings and recommendations outlined in this interim report. The TF will consider this public feedback in finalizing its recommendations for the NAIRR and in developing its roadmap and implementation plan. Importantly, this interim report does not yet offer a detailed roadmap and implementation plan for the NAIRR. Rather, following the publication of this report, the TF will continue to study and deliberate toward the generation of a more detailed plan to be released by the end of this year.

2. A Vision for the NAIRR

The NAIRR is envisioned as a large-scale, shared cyberinfrastructure that fuels AI discovery and innovation and serves a diverse set of researchers and students across a range of fields. It will help democratize access to a variety of cutting-edge computational resources by providing the data and compute capacity to support tens of thousands of users. The NAIRR will provide access to data sets and aggregate or catalog AI-relevant tools, testbeds, environments, and training resources. The NAIRR has an opportunity to both leverage and augment the Nation's existing cyberinfrastructure to advance knowledge across a variety of AI-relevant disciplines.

The TF's shared vision for the NAIRR addresses the opportunities for the U.S. AI R&D ecosystem described in [Chapter 1](#), in alignment with the goals and attributes outlined in the National AI Initiative Act of 2020. This chapter provides recommendations for the strategic objective and goals, general composition and access mechanisms, and intended user base of the NAIRR.

Strategic Objective of the NAIRR

Recommendation 2-1: The strategic objective of the NAIRR should be to strengthen and democratize the U.S. AI innovation ecosystem in a way that protects privacy, civil rights, and civil liberties.

A more robust and inclusive AI innovation ecosystem has the potential to support research whose outputs will advance U.S. economic competitiveness, improve quality of life, advance equity, strengthen national security, and advance solutions to societal challenges. This vision can be achieved by providing resources in support of research that could push the boundaries of the technology and explore new directions in AI, extend AI in novel ways across domains, advance AI testing and evaluation, leverage AI innovations to solve societal challenges, and explore AI's societal impacts.

The implementation of the NAIRR, including its long-term sustainability, must address—and will present research opportunities for—advanced cyberinfrastructure, data curation and confidentiality, security, and trustworthy AI, and thus presents an opportunity for a broad set of researchers and professionals to participate. The NAIRR should also support early experimentation by students learning how to build and apply AI so that the full breadth and diversity of talent interested in AI R&D have the opportunity to engage. Particular care should be placed on ensuring that the breadth of talent is meaningfully diverse and inclusive of traditionally underrepresented groups in AI R&D.

Goals for Establishment and Sustainment of the NAIRR

Recommendation 2-2: The NAIRR should be designed to help achieve four primary goals for AI R&D:

- (1) **Spur Innovation:** The NAIRR should support the research, development, and translation of novel methods in foundational and use-inspired AI research.

- (2) **Increase Diversity of Talent:** The NAIRR should actively seek to increase the diversity of AI researchers by lowering the barriers to participation for all, regardless of background, organizational affiliation, or geographic location within the United States.
- (3) **Improve Capacity:** The NAIRR should promote AI skills and knowledge through expanded access to AI resources, ensuring that a growing number of AI researchers in the United States are able to leverage the state of the art in their work.
- (4) **Advance Trustworthy AI:** The NAIRR should offer information, tools, and trainings in support of research that fosters the development and adoption of trustworthy and responsible AI.

Subsequent chapters of this report present more detailed findings about the current state of AI R&D and U.S. research cyberinfrastructure, TF recommendations for designing the NAIRR to achieve these goals, and approaches to evaluate its progress for achieving them.

Composition of the NAIRR

Recommendation 2-3: The NAIRR should be formulated as a federated cyberinfrastructure ecosystem run by a single management entity, with governance and external advisory bodies.

The NAIRR should broadly include the following (see also [Figure 2](#), which is explained further in [Chapter 3](#)):

- **A management entity** responsible for NAIRR operations, administration of resource allocations and access, user support, resource acquisition and partnerships, management and coordination of the user portal and other cyberinfrastructure, and more. This entity should be run by executive leadership and an operational staff.
- **A Board of Governors**, connected to the management entity, that helps to set NAIRR strategy and policies, provides long-term oversight of NAIRR operations, and to which the NAIRR executive leadership reports.
- **External advisory bodies** that provide independent oversight and evaluation of NAIRR operations as well as technical input on the composition, operations, and governance of the NAIRR.
- **Cyberinfrastructure** that comprises the federation of resources, interfaces, and expertise required to make those resources accessible as a whole. Specifically, the NAIRR cyberinfrastructure includes the following:
 - *Resources:* The set of compute, data, edge, and other capabilities that could be owned and operated by governmental, non-governmental, or private sector entities or by the NAIRR management entity itself. For the purposes of this report, resources include the cyberinfrastructure enabling access to AI testbeds, as well as the testbeds themselves.

- *Interface*: A central, user-facing but non-exclusive interface to NAIRR resources that also permits other pathways to resource access. This portal will also provide catalogs and associated search and discovery tools for easy access to educational and training materials to support the use of the NAIRR in educational contexts.
- *Expertise*: The expertise necessary to manage and support the resources and interfaces of the NAIRR, including providing training and support to NAIRR users.
- **Education/training resources**, including the technical training and support related to the use of the NAIRR, tailored to a range of experience levels.

The NAIRR should be built by knitting together both existing and new resources to form a cohesive, accessible cyberinfrastructure based on common standards and clear governance processes. The system should take advantage of existing campus-, regional-, and national-scale resources when possible but also add significant new resources to meet the national need. Additionally, it should connect users to a diverse set of public- and private-sector data, compute, testbed, and other resources through an intuitive interface, facilitating their use through educational tools and user support. Importantly, NAIRR computational resources should span the full range of possible offerings, including commercial cloud, high-performance and high-throughput computing, on-premise (at academic and/or government sites) resources, "edge" computing resources and devices, and novel computing approaches and platforms. The NAIRR should be designed to embody attributes of transparency and trust, security and robustness, accessibility, scalable functionality, sustainability, and accountability.

Access to the NAIRR

Recommendation 2-4: NAIRR user access policies should be grounded in the principles of equity, fairness, and security.

Detailed access policies and a plan for upholding these principles should be defined with consideration of stakeholder input and published transparently as part of NAIRR governance policies.

Recommendation 2-5: The NAIRR management entity should work toward providing an integrated access portal through which all resources are made available.

This integrated access portal would not restrict federated resource providers who affiliate with the NAIRR from continuing independently to make their resources available to users. More information on the NAIRR user interface can be found in [Chapter 4](#).

NAIRR User Base

Recommendation 2-6: The NAIRR should be designed to support the needs of researchers and students from diverse backgrounds who are pursuing foundational, use-inspired, and translational AI research.

The NAIRR should support the full spectrum of AI research—from foundational to use-inspired to translational (refer to [Appendix A](#) for relevant definitions). The three primary user profiles for the NAIRR are envisioned as follows:

1. *AI researchers*: those who are advancing the state of the art in AI and developing approaches in pursuit of improved and novel capabilities leveraging AI innovations;
2. *Researchers conducting use-inspired AI research and using AI innovations to advance other fields*: those who are advancing AI or developing innovative applications of AI to solve problems in another domain of study, including science, engineering, medicine, business, and the humanities (while also furthering AI itself); and
3. *Students*: those studying at community colleges, four-year colleges and universities, or graduate schools who are learning about AI and experimenting with the development of AI models and tools.

The primary users of the NAIRR would be U.S.-based AI researchers and students who fit the user profiles detailed above and are affiliated with U.S. organizations of the following types:

- Academic institutions;
- Non-profit organizations;
- Federal agencies or federally funded research and development centers (FFRDCs); or
- Startups or small businesses 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, to advance foundational, use-inspired, or translational AI R&D.¹⁸

Private sector researchers other than small businesses with Federal funding would be allowed to access NAIRR resources, but only at limited levels and in support of research that is in the public interest, as determined via evaluation of private sector research proposals, subject to NAIRR policies and user agreements, and with the requirement that results from this work be made publicly available.

Access to NAIRR resources for individuals within this user base would be subject to an evaluation process as defined by the NAIRR management entity. (See [Chapter 3](#) for additional information on oversight and resource allocation.)

3. Establishing and Sustaining the NAIRR

In order to establish a NAIRR capable of achieving the goals outlined in [Chapter 2](#) that is responsive to user needs and has a high probability of success, the appropriate people and organizations must be engaged. To be successful, the NAIRR will require funding, organizational and operational management, a governance structure, and oversight mechanisms. This chapter provides preliminary recommendations in these areas.

Agency Funding, Roles, and Responsibilities

The U.S. Federal Government plays an important role in the U.S. innovation ecosystem by funding foundational, use-inspired, and translational research and research infrastructure. Federal agencies currently fund a range of AI-related research and support a variety of research cyberinfrastructure resources that are hosted and managed by universities and national labs, and are in turn made available across a range of research communities. The NAIRR would transform the U.S. national research cyberinfrastructure with its unique mandate of strengthening and democratizing foundational, use-inspired, and translational AI R&D in the Nation. The objectives of the NAIRR align with the missions of several Federal science agencies, including the Department of Energy (DOE), NSF, National Institutes of Health (NIH), and National Institute of Standards and Technology (NIST). Indeed, by embracing a federated approach, the NAIRR could help agencies better leverage their collective investments and build future collaborations.

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.

As no single agency owns the mission space of advancing the Nation's AI capability, and as the NAIRR will inherently require investment in new resources in addition to the aggregation of existing capabilities, Congress should fund the NAIRR through appropriations to a collection of Federal agencies that represent AI stakeholders. The NAIRR management entity could also explore approaches that contribute to its long-term sustainability through other funding or revenue sources, consistent with NAIRR goals and principles and in compliance with legal and regulatory requirements.

Recommendation 3-2: Representatives from each of the sponsoring agencies should engage with NAIRR management and administration to provide expertise and oversight.

These agencies will have a vested interest in the success of the NAIRR and can provide expertise and oversight. They will also play an important role in supporting community building and identifying stakeholder needs among their research communities, and in maintaining situational awareness to inform decisions about complementary Federal programs and investments. It will be incumbent on the NAIRR management entity to integrate the perspectives of various Federal stakeholders into a cohesive strategy.

Recommendation 3-3: Federal agencies should make relevant cyberinfrastructure resources available via the NAIRR.

Agencies should make new and existing federally funded or federally owned cyberinfrastructure, such as data sets, computational resources, software and services, and testbeds, available via the NAIRR, expanding scope and scale as needed and when resources and funding permit. Achieving the vision for the NAIRR will likely require NAIRR-specific funding to expand the capacity of existing agency resources. In the case of sensitive resources, specific decisions about which resources to include will depend on legal and ethical considerations and on the robustness of tools and policies for protecting confidentiality, privacy, and civil rights and civil liberties (see [Chapters 4, 5, and 6](#) for more details).

Ownership and Administration

Federally funded research cyberinfrastructure is currently managed via a range of different models, each having evolved to meet the mission needs of respective agencies. The NAIRR has a distinct mission, which—although supported by many agencies—is entirely unique: to expand and democratize access to AI R&D resources broadly across the Nation. While management of such an infrastructure could be handled directly by forming a new division or element under an existing government agency (which has the benefit of assigning clear ownership, authority, and responsibility), assigning a single agency to manage the NAIRR could risk narrowing its focus to that agency's specific mission, leaving the needs of researchers supported by other agencies unmet. Furthermore, agencies often face cumbersome contracting and procurement requirements that can result in long lead times and inhibit agility in managing such a resource.

There are several alternatives. One option would be for multiple Federal agencies to come together to manage the NAIRR through an appropriately authorized and funded interagency coordination office. Such an approach would require, among other things, new authorities to hire appropriate talent and manage resource acquisition. An alternative option would be for a university to manage the NAIRR in a manner similar to NSF-sponsored advanced cyberinfrastructure resources. Universities generally have the advantage of established administrative infrastructure for grant application and management, and inherently focus on students and academic researchers, who are central to the intended NAIRR user base. A third option would be a consortium model,¹⁹ in which multiple stakeholders come together around a shared goal. This option leverages strengths of each contributor and tends to offer agility—although conflicts of interest and intellectual property issues can be challenging to navigate across multiple stakeholders. A fourth option is to leverage FFRDCs, which provide the Federal Government with sustainable and persistent capacity to address long-term problems that cannot be addressed as effectively using existing in-house government or contractor resources. FFRDCs must be operated, managed, or administered by an autonomous organization in line with the public interest and free from conflict of interest,²⁰ can have special access to government data,²¹ and can report to multiple agencies by design and thus address the needs of multiple stakeholders. Other options may also exist.

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

The NAIRR management entity could be an FFRDC, a university, a contractor, a non-profit organization, an institute, a consortium, or another such entity. Although Federal oversight is anticipated regardless of the approach employed, in considering the different models above, it is important to assess the relative risk factors associated with data security, rights, and privacy for each approach. Regardless of the approach, the NAIRR management entity will require a permanent and diverse staff focused on resource provisioning, managing core operations and system components—including cataloging external resources, providing user support, and overseeing security operations (as indicated in [Figure 2](#), which is explained further in the Governance and Oversight section of this chapter). Ideally, such staff would be primarily dedicated to delivering on the NAIRR mission.

The executive leadership of the NAIRR—for example, a Director, Chief Executive Officer, or management team—would make day-to-day operational decisions, oversee and manage NAIRR staff, report to the board of governors, and conduct outreach and communication activities to external stakeholders, such as evaluators, advisors, users, and the public, as well as engage potential partners to participate in the NAIRR. With input from its advisory bodies, NAIRR leadership should also actively scout both (1) new cyberinfrastructure resources and capabilities as technologies advance and (2) emerging trends and needs in the AI R&D community, and update NAIRR resources and services accordingly.




Figure 2. Key functional roles for NAIRR management, governance, operations, and oversight

Recommendation 3-5: The NAIRR should collaborate with resource providers to make a broad variety of resources available through the NAIRR user access portal.

While new cyberinfrastructure resources may be required to serve NAIRR users, the NAIRR need not duplicate systems that are already available and positioned to support NAIRR user needs. Instead, the NAIRR should aim to leverage existing resources through appropriate partnerships or other agreements if another entity is better positioned to provide them. The executive leadership would be responsible for nurturing such partnerships, and NAIRR staff would manage resource linkage and system interoperability, working in close collaboration with resource providers.

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.

Similarly, the NAIRR management entity may choose to but need not design, build, or support a particular component of the cyberinfrastructure if other entities are better positioned to do this efficiently, at a lower cost, and in alignment with the vision and guiding principles of the NAIRR. Resource owners should have primary responsibility for managing their systems in compliance with transparent NAIRR policies and standards. Other NAIRR operational responsibilities would be distributed among the NAIRR management entity, federated resource providers, and possibly contractors via public-private partnerships, or other agreement types, according to the NAIRR management entity's operations plan.

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.

Because the NAIRR will be a newly established entity, there is an opportunity to build a DEIA focus into the system and operational plan from the beginning, rather than as an afterthought. Extending access to AI research resources as broadly as possible is a fundamental goal of the NAIRR. Assessment of progress and input on engagement and support of a broad and diverse AI community will be a key aspect of NAIRR governance and oversight activities.

Governance and Oversight

The NAIRR is envisioned as an operational resource that provides access for tens of thousands of users.²² Ensuring that the system fulfills its vision will require strategic planning, policies, and system design, as well as ongoing oversight, evaluation of the system's performance, and adaptation to meet NAIRR user needs. Such activities can be assigned to designated entities, with membership including some diverse combination of NAIRR staff and outside experts, spanning a range of demographics and relevant expertise, experience, and backgrounds, depending on the associated functions and the frequency and intensity of effort involved.

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.

The NAIRR should operate under guiding principles of transparency and trust, and provide clear documentation of all policies, practices, and expectations of users, contractors, and partner entities. These policies will need to define security controls, access rights and responsibilities, codes of conduct, minimum privacy and confidentiality protections, technical policies and standards, an ethical framework for NAIRR use, oversight and accountability processes, and a business model and financial plan, among others. These policies should ensure legal and regulatory compliance and be designed through consultations with the user community to support the strategic objective and goals of the NAIRR. They should also define composition, roles, and terms of service of individuals who serve in a governance role. Finally, the charter or other governance documents should also identify key performance metrics and indicators that can be used to assess performance and impact of the 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.

As shown in [Figure 2](#), a potential set of organizational roles would include executive leadership (such as a NAIRR Director or Chief Executive Officer), accountable to a board of governors broadly representing the research interests of multiple Federal stakeholders as well as academia and the private sector. In addition, a set of advisory, oversight, or evaluation boards (composed of outside experts) should convene periodically to evaluate overall performance against metrics defined in the NAIRR charter or governance documents and advise on a variety of technical areas, as indicated in [Figure 2](#). For example, the scientific merit of research projects completed using the NAIRR should be evaluated to assess the impact of the NAIRR on the field of AI and on other science and engineering fields that use AI. The NAIRR governance structure should be designed to provide oversight on the ethical use of NAIRR resources; for example, by defining policies to protect privacy, civil rights, and civil liberties, and by screening NAIRR projects or proposals for potential harms. The NAIRR charter should also provide a mechanism for renewing or sunseting the activity as indicated by its assessed performance and the evolving needs of the AI R&D ecosystem (see also Recommendation 3-20).

As part of the core operations of the NAIRR, ongoing research projects and new research proposals supported by NAIRR resources will need to be reviewed at regular intervals for ethics policy compliance. Project management staff should oversee research proposal review, administer resource (e.g., compute and data) allocation and access processes, provide user support, and serve as an interface to NAIRR user communities. Resource management staff should oversee and run NAIRR-hosted infrastructure and services and serve as interfaces to the resource provider community partnering with the NAIRR. NAIRR staff and governance mechanisms should also monitor and facilitate partner entities' compliance with NAIRR policies and standards.

Resource Allocation and Sustainment

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.

All users of NAIRR computational resources should be required to pass a research proposal evaluation process, administered either through a Federal funding agency or through the NAIRR management entity including the board of governors. Outputs of research, including intellectual property and potential patents, resulting from access to the NAIRR will be subject to existing policies of the Federal agencies funding that research, or similar policies established by a NAIRR management entity for research supported by other means. For example, many Federal agencies require that peer-reviewed publications or juried conference proceedings resulting from the research they fund be deposited in public-access compliant repositories within 12 months after initial publication.²³ They also have requirements for data management and sharing that aim to maximize the availability of data generated by funded research.

Access to all resources, including data, should be governed by clear use policies and user agreements to which users are held accountable and that the system design reinforces. The NAIRR management entity should aim to implement standard legal agreements for users and resource providers, establishing common terms of use and clarifying that users retain the intellectual property they develop using NAIRR resources. Such 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.

Recommendation 3-11: NAIRR resource allocation processes should provide accessible on-ramps for the range of expected users, be lightweight where possible, and be as inclusive as feasible.

Resources should be allocated such that (1) researchers who receive Federal funding can access resources with minimal additional administrative burden (e.g., no additional proposals) and (2) researchers and students without current Federal support can apply for access. For small, startup allocations, the NAIRR management entity should manage a rapid review process to expedite user access. Special attention should be given to underserved and underrepresented research communities throughout the Nation in alignment with the goal of the NAIRR of increasing diversity among AI researchers.

Recommendation 3-12: Costs of allocations should follow a tiered model, in which some resources are fee-based and some are provided at no cost.

NAIRR management and governance should define a transparent fee structure consistent with applicable Federal laws and regulations, ideally with some access heavily subsidized or free. The lowest cost resources should be made more widely accessible to the broadest range of users. Researchers should, in general, be able to apply and compete for allocations of limited resources reserved in a dedicated pool, with decisions informed by research merit and the objective of broadening access and participation. Researchers who have received grants from sponsoring agencies could receive allocations, or funding toward allocations, as part of their award.

Recommendation 3-13: The framework for NAIRR resource allocation should be designed to incentivize contributions to the NAIRR user community or to the public good.

Mechanisms to incentivize contributions to the NAIRR would help to strengthen the overall NAIRR ecosystem and support its sustainment. A "leaderboard" approach, whereby users may compete for recognition as top contributors to the community good, could also be considered.

Recommendation 3-14: The NAIRR management entity should also explore mechanisms to incentivize data and metadata contributions that would help to add value to the resource.

As an example, private entities wanting the research community to extract insights from their own data could make these data available to researchers through the NAIRR platform. Researchers would benefit from access to these resources, and the data owner could glean insights from the associated research. This model would incentivize data owners to clean, label, and validate their own data and would enhance usability of the data for AI R&D.

NAIRR Performance Indicators and Metrics

As the NAIRR will be a national and publicly funded resource, it will be important for the NAIRR management entity to document and disseminate the value of the research it supports, along with its broader impacts. In particular, the NAIRR management entity should provide such documentation to the American public, Congress, and the sponsoring Federal agencies in an open, reproducible, and transparent manner.

Determining the value and impact of research programs is nontrivial but is facilitated by advance planning. Designing a program strategically to target clearly defined and measurable goals at the outset also increases its probability of success. Periodic program evaluation and collection of stakeholder input can inform decision making and enable adjustment as needs and opportunities change.

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

Successful program design requires a clear understanding of the activity's overall near- and long-term goals and a conceptual model for how progress toward achieving them will be measured. Such a so-called logic model defines the underlying assumptions of how the program would work, its desired outcomes and impacts, the activities and outputs that would support these outcomes, the resources or inputs needed to achieve the outcomes, how all these elements work together to achieve desired impacts, and the metrics that can be used to characterize each.²⁴ The entities responsible for evaluating the NAIRR should adopt this approach to assessing its performance, similarly informing the evolution of the NAIRR over time.

Recommendation 3-16: Qualified external evaluators should regularly assess the performance of the NAIRR.

Independent, external evaluators should assess the performance of the NAIRR using rigorous and transparent methods, and disseminate their findings to the public, Congress, and the supporting Federal agencies. Evaluation data on inputs (e.g., funding, staffing, and other resources) and activities (e.g., NAIRR use, research efforts, outreach, and creation of reference materials and catalogs) should be drawn from a range of sources, including NAIRR data.

The external evaluators, in conjunction with the NAIRR management entity and the NAIRR board of governors, should use established best practices to evaluate each set of programmatic investments once the scale of these investments is determined to gauge the value of the role of the NAIRR in advancing the AI research ecosystem. These measures should be complemented by annual user surveys and ongoing user engagement mechanisms to gauge user satisfaction, assess platform usability, identify emerging needs, and capture researcher-level outcomes.

Recommendation 3-17: The NAIRR management entity should capture and record data to support evaluation of the resource, and such data should be shared with the public in a timely manner.

Data should be captured to support four levels of performance indicators: (1) measures associated with resource investments; (2) measures of resource usage/activities, including user diversity; (3) measures of outputs; and (4) measures of impact of the resource and the research it enables. Such data should be shared with the public in a timely manner, subject to reasonable restrictions to protect privacy and security.

The indicators should be regularly assessed to gauge progress toward intended outcomes and inform NAIRR resource investment and service management decisions, driving a nimble approach that is responsive to the user community. It may be necessary to adopt new indicators over time as the NAIRR evolves.

Recommendation 3-18: The NAIRR management entity should budget sufficient funds for robust data collection and evaluation activities, and design NAIRR processes to capture key analytics measures from the outset.

Efforts to collect baseline data around measures of the composition, strength, and productivity of the AI research community should begin as soon as feasible, in parallel with the development and launch of the NAIRR.

The NAIRR management entity should additionally leverage national data sources, such as NSF's National Center for Science and Engineering Statistics, state-level administrative records on education and workforce, institutional human resources and financial records (when already made available), and ongoing annual surveys,²⁵ to track national progress toward the more indirect, strategic impact goals of the 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.

To facilitate this tracking, the NAIRR management entity should provide standardized acknowledgement language that all NAIRR users are required to use when disseminating the results of their research, including data, models, tools, publications, and patents. This standardization should be complemented by requesting that users cite a "marker" publication on the NAIRR, facilitating automated retrieval of relevant research outputs via existing databases and search engines for scholarly publications.

Recommendation 3-20: The NAIRR charter should establish mechanisms by which evaluation can directly inform adjustments to the strategic goals, operational functions, and resource capabilities of the NAIRR to optimize impact and meet changing needs.

Evaluation is important not only for providing oversight and transparency but also for informing improvements to the activity. Results of assessment of both NAIRR performance and changing user and research community needs should feed into adjustments to NAIRR goals, resources, governance, and operations. This assessment would also include a mechanism for sunseting the NAIRR should a different model be deemed necessary as the field and user needs change, or if the activity is found to be unsuccessful. (See also Recommendation 3-9.)

4. NAIRR Resource Elements and Capabilities

The National AI Initiative Act of 2020 directs the NAIRR roadmap and implementation plan to include findings and recommendations on the "capabilities required to create and maintain a shared computing infrastructure to facilitate access to computing resources for researchers across the country, including scalability, secured access control, resident data engineering and curation expertise, provision of curated data sets, computational resources, educational tools and services, and a user interface portal."

This chapter presents the TF's preliminary findings about AI R&D resources and needs, along with recommendations for how the NAIRR, once established, could provide access to an integrated set of resources for the AI R&D community. It also describes desirable attributes of these resources and some of the necessary capabilities for achieving them. In addition to data, compute, and educational resources, there are opportunities for the NAIRR to provide access to testbeds and other testing tools for foundational, use-inspired, and translational AI R&D.

Data

Progress in many areas of AI has been made possible by the increasing availability of data with which to train and test data-driven AI models and has been enabled by the ongoing digitization of scientific, business, and social practices. Yet important data resources are not equally available to all who wish to advance understanding and capabilities in AI, as described in [Chapter 1](#).

The TF considered the current research data landscape; potential sources of AI-ready data; and ethical, privacy, security, and usability concerns with different use cases, data types, and models in the context of the strategic objective and goals of the NAIRR. This section provides findings about data needs for AI R&D along with preliminary recommendations for features of the NAIRR that would facilitate responsible access to critical data resources for a broad base of researchers.

Findings

Finding 4-1: Rigorous AI R&D is often not possible without high-quality, trusted, dense, and transparent data resources.

The performance of many AI models depends on the quality of the data on which the models were trained and/or tested. If data are inaccurate or biased, the models could be more likely to propagate bias or yield inaccurate or harmful results, depending on how the models are deployed. In addition, data sets that are not dense—that is, they do not include a significant number of examples of a feature or phenomenon of interest (sparse data)—may be insufficient to develop a rigorous model. Researchers who lack access to such rigorous data will not be competitive at the leading edge, and risk developing AI models that could lead to biased, inaccurate, or insecure systems.

Finding 4-2: There are substantial data quality challenges within and across most research domains.

In recent years, science and industry have become increasingly data driven. Data are often highly distributed, not discoverable, and insufficiently reusable. In general, data are also extremely heterogeneous (e.g., video, voice, text, image, sensor) and poorly documented and curated. A wide variety of data types are of potential interest for AI R&D. Many data sets are unstructured and unlabeled or have few labeled examples. While many approaches to AI research require structured and labeled data, other AI research can make use of any type of available data.

Finding 4-3: Data curation is a substantial challenge for researchers in all domains.

Labeling, tagging, and annotation are difficult to automate and require significant hours of expert analysis. Data labeling and curation standards are generally evolving, limited, absent, or inconsistently adopted throughout research communities. Research data curation generally relies on communities of expert researchers from academia or the private sector. In some areas, standards are ad hoc or are not established or not adopted.

Finding 4-4: There are substantial costs to combining and linking heterogeneous data.

Building AI-ready data sets entails an ongoing process of care, including curating, labeling, maintaining, transferring, and retiring data sets, which requires expert human effort and substantial time.²⁶ Transfer of large data sets also requires time and effort, and can be expensive when leveraging commercial cloud resources.

The use of many types of data sets for R&D (e.g., data pertaining to real people) poses valid privacy concerns, but mitigation options may exist that are deemed sufficient to manage the associated risks of an intended use. In general, data privacy risks must be considered and managed in the context of the potential value of the research that would leverage the data in question.

Finding 4-5: There are opportunities to learn from, complement, or build on ongoing activities to leverage data for the public good.

The implementation of the Foundations for Evidence-Based Policymaking Act is addressing statistical and administrative data access challenges, including access to restricted data.²⁷ In particular, the Advisory Committee on Data for Evidence Building has worked closely with Federal statistical and programmatic agencies, as well as state and local governments, to facilitate data linkage and sharing, promote data access, and demonstrate the societal value of leveraging existing data resources.²⁸ In addition, the OSTP Memorandum on Increasing Access to the Results of Federally Funded Scientific Research is prompting Federal science agencies to make more research data accessible to the public, and the recommendations from the Equitable Data Working Group lay out a future for Federal data sharing with the research community in a manner that builds capacity for equitable policymaking and outcomes.²⁹

Recommendations

The NAIRR will require investment in several infrastructure elements. These include a technical infrastructure to host data in secure facilities so that costs are minimized and secure access is maximized; an access infrastructure that is networked to enable the domain-specific heterogeneity of data structures to be addressed; trained NAIRR staff who support diverse community data curation, linkage, and validation activities; training programs to develop a diverse AI workforce, foster innovation, and create community-driven value in the context of the NAIRR; and a search and discovery platform so that knowledge about NAIRR data use, users, and value can be identified, leveraged, and replicated.

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.

The sheer volume and variety of data of interest will make it impossible for the NAIRR to curate any or all of it; the ability to leverage important resource elements managed by partner entities as well as promote appropriate best practices is a key feature of a federated system. However, the core NAIRR management entity should also contribute by facilitating the curation of data that can be made widely available to users. Data resources could be contributed by researchers; non-profit or commercial organizations; government agencies; state, local, and/or tribal governments; academic institutions; and citizen scientists.

Periodic review of the quality and breadth of data hosted by or made available through the NAIRR, including alignment with best practices, compliance with ethics policies and the Findability, Accessibility, Interoperability, and Reusability (FAIR) principles, is also needed.³⁰

Recommendation 4-2: Privacy should be protected by following the "Five Safes" framework for safe use (safe projects, safe people, safe data, safe settings, and safe outputs).

This is also a key attribute for NAIRR security (see [Chapter 5](#) for additional details).

Recommendation 4-3: NAIRR leadership should establish and periodically update policies and governance structures that address, measure, and report on data quality, use, and equitable access.

(See [Chapter 3](#) for additional details.)

Recommendation 4-4: The NAIRR should protect civil liberties and civil rights by using best-in-class technical approaches to establish governance policies, fund oversight entities and evaluations, enable public engagement, promote transparency, and reduce bias and potential harms associated with NAIRR data use.

(See [Chapter 6](#) for additional details.)

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.

This approach could be achieved by demonstrating the value generated in the form of better AI R&D tools and stronger AI models that facilitate scientific discovery or address real-world problems; recognizing individual data providers or curators whose contributions have had noteworthy impact; returning AI-driven insights to data holders who contribute data with research value associated with a particular domain; and building a collaborative community where contributions are rewarded in kind. (See Recommendations 3-13 and 3-14 for additional details.)

Recommendation 4-6: The NAIRR should provide infrastructure and permanent staff support for provider, host, or community data curation and incentivize the continued development of community-driven standards and improvements to data quality as determined by the relevant domains, in consultation with NAIRR staff.

Substantial resources should be dedicated to technical support staff. Data users, contributors, and curators will require support to understand and meet technical standards and to ensure rigorous data use. The growing heterogeneity and complexity of data mean that substantial knowledge of the system will be required from support staff. Attracting and retaining skilled staff with appropriate compensation packages will be essential.

Recommendation 4-7: Substantial resources should be devoted to the establishment of training programs on NAIRR data policies, use, and curation to maximize the capacity of domain communities to provide substantive contributions and to promote best security practices.

Such training offerings will be critical to the confidentiality, integrity, and assurance of the NAIRR, and to protect privacy, civil rights, and civil liberties. This recommendation is consistent with both the Foundations for Evidence-Based Policymaking Act and the Federal Data Strategy. (See [Chapters 5](#) and [6](#) for additional details.)

Recommendation 4-8: The NAIRR should establish a value ecosystem around data that can be used for AI, and support data search and discovery.

The NAIRR should provide high-value, core data sets to establish a value proposition and jump-start search and discovery. It should use digital tools to identify key data sets for different research areas; partner with academic and local communities to bring in additional data and knowledge; and engage governors, legislators, and local decision makers to ensure that the data resources provide value to the research community and benefits to the public.

Recommendation 4-9: If the decision is made for the NAIRR to include confidential or sensitive data, access to such data should be tiered, controlled by the data providers, and available through the same portal through which compute resources are provided.

In addition to widely accessible, open data sets suitable for broad use, if the NAIRR provides access to confidential or sensitive data sets, such access should be provided to users according to need and potential benefits of their work, and with appropriate access control requirements including controls by data providers (see Government Data Sets and User Interface in [Chapter 4](#), and [Chapters 5](#) and [6](#) for additional details).

Government Data Sets

Government agencies hold data that could fuel foundational, use-inspired, and translational AI research. Agencies' data collection, reporting, and sharing requirements, such as those associated with the Foundations for Evidence-Based Policymaking Act, the OSTP Memorandum on Increasing Access to the Results of Federally Funded Scientific Research, and the recommendations from the Equitable Data Working Group could be leveraged to yield data contributions to the NAIRR. Key domains in which the Federal Government could help drive AI-based innovation are transportation, healthcare, and natural hazards research, among many others. Each category has different characteristics, requiring different approaches for how the corresponding data might be integrated into the NAIRR and the technical and policy protections required. Sources of Federal agency data include statistical data, administrative data, and data from federally funded intramural and extramural research; these could be highly valuable to AI research, and the NAIRR could help make the data accessible.

Federally held data generally come with acquisition, curation, distribution, and management challenges rooted in budgetary constraints, staffing limitations, legal requirements, lack of cohesive data access mechanisms or interoperability standards, low data quality, and coordination of incentives.³¹ In addition, agencies often have legal responsibilities that prevent or create duties related to data sharing. For example, the use of Federal statistical data is subject to the Confidential Information Protection and Statistical Efficiency Act. Census data are subject to U.S.C. Title 13, and Federal tax information use is subject to U.S.C. Title 26. Agencies need staff to document, share, manage, and create standards for data, and funding to support data sharing efforts across departments.³² The 2021 Office of Personnel Management designation of Data Scientist as a new occupational series is a step toward growing the necessary human capital for this effort.³³

Users requesting data from Federal agencies face challenges managing the heterogeneity, complexity, and timelines of data sharing agreements. Data sharing arrangements vary across agencies: Some have dedicated offices while others use ad hoc processes, and data sharing agreements are often only one-off arrangements and rely on interpersonal connections.³⁴ Moreover, the burden of navigating legal and regulatory hurdles often falls on the researcher. The negotiation of data use agreements (DUAs) for university researchers to use agency data can be complex, resource intensive, and lengthy. Current barriers to the execution of DUAs include a lack of sufficient university contract negotiators, lack of incentives for the data provider to grant access to the data, and challenges negotiating contractual provisions around indemnification. Information privacy and security protections, ownership of university researchers' work products, and

appropriate mechanisms for dissemination of results must also be established. Universities need support navigating the DUA process. Standardizing the DUA process could greatly decrease the burden to data requesters and providers.³⁵

The TF considered current barriers to leveraging government-held data for R&D and developed several preliminary recommendations for how the NAIRR could approach incorporation of Federal data to improve the ability of researchers to leverage it for accepted uses.

Recommendations

Recommendation 4-10: The NAIRR management entity should explore making three types of government data available through the NAIRR: statistical data, administrative data, and data generated by federally funded research.

The Federal Statistical System provides an important element of the data collection, surveys, processing, and analysis needed for evidence-based policymaking, such as the Decennial Census, and is composed of 13 principal statistical agencies and 3 statistical units with extensive statistical data.³⁶

Another important element of Federal data collection results from the administration of government programs. These regularly collected data could be repurposed for AI R&D. Programmatic and transaction data could come from sources such as the U.S. Department of Health and Human Services,³⁷ United States Department of Veterans Affairs, U.S. Department of Defense, Department of Justice, Department of Transportation, U.S. Department of Agriculture, National Oceanic and Atmospheric Administration (NOAA), Department of Homeland Security (DHS), National Aeronautics and Space Administration (NASA), and Department of Housing and Urban Development (HUD).

Finally, data generated by federally funded research are a rich source for foundational, use-inspired, and translational AI research.³⁸ Federally held scientific data includes results of research funded by agencies such as DHS, NASA, NOAA, NSF, and NIH. For example, the NAIRR could help make accessible satellite data collected by NASA on changing ice, cloud, and land elevations to advance AI-enabled research on climate change.³⁹

Recommendation 4-11: The NAIRR management entity should build on and leverage existing Federal data sharing efforts to facilitate access to data sets to researchers for approved uses, with its infrastructure designed to support the necessary controls and protections.

NAIRR priorities may align with those outlined in the Foundations for Evidence-Based Policymaking Act of 2018, the Geospatial Data Act of 2018, Executive Order 13859 on Maintaining American Leadership in Artificial Intelligence, the Office of Management and Budget (OMB) memorandum on Improving Implementation of the Information Quality Act, the OSTP Memorandum on Increasing Access to the Results of Federally Funded Scientific Research, the 2021 Federal Data Strategy Action Plan, and the recommendations from the Equitable Data Working Group. Advancing the goals of the 2020 Federal Data Strategy Action Plan, the 2021 Federal Data Strategy Action Plan directs agencies to identify and inventory priority data assets,

including assets in support of AI R&D, for Open Data Plans.⁴⁰ The NAIRR management entity should identify synergies or coordinate with the efforts of OMB's Office of Information and Regulatory Affairs and Office of E-Government and Information Technology,⁴¹ and the Chief Data Officer Council⁴² established by the Foundations for Evidence-Based Policymaking Act.

Compute Resources

The growing utility and prominence of AI-based learning systems, while enabled by data availability, is possible only with the parallel advances in processing capabilities and storage capacity of computer hardware. Both high-quality data and substantial compute capabilities are now necessary for research in a growing range of AI domains.

The NAIRR has an imperative to lower barriers to entry into AI research by providing access to computational and data resources for a variety of new users who otherwise would face challenges contributing. Leveraging existing resources may accelerate time to launch and enable flexibility in providing a variety of state-of-the-art tools needed by the user community; however, many existing resources are oversubscribed or are operating at or near capacity. Expanding access to computing resources will require adding substantial new research computing infrastructure and capabilities.

This section provides preliminary findings and recommendations for the design of a federated compute research infrastructure that supports foundational, use-inspired, and translational AI research across a broad user base by providing both cutting-edge capabilities (i.e., to run the largest research problems in AI) and the capacity to provide resources to a large number of users. Given that technologies evolve rapidly, the TF makes no vendor- or product-specific recommendations, but rather outlines some general principles the NAIRR management entity should adhere to when acquiring computing resources or identifying partner resource providers.

Findings

Finding 4-6: AI R&D leverages both "production" and "experimental" compute resources.

A "production" system can be used to run state-of-the-art software tools, frameworks, and applications without modification and with reasonable expectations of stability, reliability, and quality of service. For example, standard frameworks most researchers use today, such as PyTorch⁴³ and TensorFlow,⁴⁴ should run smoothly and not require substantial porting and optimization efforts to use them. An "experimental" system or resource that is exploring a new hardware or software capability may provide an immature or rapidly evolving environment to the user. In this sort of computational environment, users may need to undertake additional efforts to port applications to properly use the capabilities of the system, rather than the "turnkey" environment expected from production systems. Some use cases may not be well supported on an experimental system. Hardware that changes the programming paradigm substantively would be considered an experimental resource.

Finding 4-7: Computing technologies, from hardware to software, evolve rapidly.

It is difficult to predict which hardware and software will be competitive or desirable in the future. Thus, the service vendors or products determined to best support specific AI needs may also change over time.

Finding 4-8: Within a research cyberinfrastructure that handles high-volume data, effective computing may require the co-location of data with the hardware on which it will be processed.

AI training data sets can be many terabytes (TB) in size. For example, MLCommons.org offers two published speech recognition data sets that are 0.24 TB and 1 TB in size.⁴⁵ With current technology, moving 1 TB of data over the commercial internet would require many hours at today's prevailing speeds. High-performance regional optical networks operate in the 1 gigabit per second (Gbps) range, reducing the time to several hours; that time is still relatively long. Larger data sets would take proportionately longer. High-performance fiber optic local area networks inside a computer center operate in the 10–100 Gbps range, dropping the data access time from 10 minutes to 1 minute; that time is tolerable. Data of 1 TB residing on the parallel file system of an HPC cluster is accessed at 400–1000 Gbps rates in just a few seconds, a small fraction of the time required to train the model. Ultimately, the closer the data is to the computing resources, the more efficiently the system can operate.

Finding 4-9: Edge computing is an emerging area of promise for AI R&D.

"Edge computing" refers to the placement of computing capability on devices connected at the edges of a network, enabling decentralized data processing, highly scalable processing capacities, and more adaptive, responsive, and accurate real-time controls. It is estimated that, by 2025, over 30 billion Internet of Things/connected edge devices will be creating approximately 4.6 trillion gigabytes of data per day.⁴⁶ This calculation means that the number of connected edge devices globally will more than double, and the amount of data created per day will more than quadruple over the next 3 years. Core technologies and standards to support efficient and safe operation of decentralized edge computing are currently lacking.

Recommendations

Recommendation 4-12: The NAIRR should provide a diverse set of stakeholders with access to a federated mix of on-premise and commercial computational resources, including conventional servers, computing clusters, HPC, and cloud computing.

The federation should include both existing resources, acknowledging that many are heavily subscribed and may require investment to expand their capacity, and new resources procured specifically for the NAIRR. This mix of resources will provide the variety needed to address a range of user needs for conducting foundational, use-inspired, and translational AI R&D. A federation of Government-procured (academic or Government lab) resources—for example, analogous to the eXtreme Science and Engineering Discovery Environment (XSEDE) Service Provider model⁴⁷—could provide for large single-capability runs, additional heterogeneity, and innovation, while the commercial cloud could be used for scalable capacity runs. This resource

mix should be developed with capacity, capability, and cost effectiveness in mind. To the extent possible, these federated resources should be co-located with traditional scientific computing resources and data resources and also include sufficient network capacity.

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

System software is dynamic and must evolve in the face of changing heterogeneous hardware. While this requirement makes it difficult to forecast future needs, three general levels can be considered. The first level is the infrastructure building block level, in which users can compose compute, storage, and network elements into a private virtual machine/cluster customized to their needs. This level would appeal to expert users accustomed to building their own on-premise systems (e.g., computer scientists). The second level is the AI application level, in which standard software packages, including TensorFlow,⁴⁸ PyTorch,⁴⁹ and Apache Spark,⁵⁰ are pre-installed and available for use on suitably powerful computing infrastructures. This level would appeal to a broad cross section of use-inspired AI researchers who want to use familiar packages in the NAIRR. At the third and highest level, the NAIRR could provide access to "serverless" AI applications, or application programming interfaces (APIs), to services provided by the NAIRR or NAIRR-affiliated organizations. Examples could include natural language processing, optical character recognition, knowledge graphs, scientific tools such as BLAST,⁵¹ as well as standard data source access (e.g., Sky Search from an astronomy image repository). This level would appeal to novice users as well as experienced domain scientists who access compute and data resources through highly structured domain-specific portals.

Recommendation 4-14: The NAIRR should not house production applications that can reduce overall capacity, but should enable services supporting AI R&D to be piloted and scaled as appropriate.

Innovative AI applications can be prototyped in the NAIRR but should find permanent homes elsewhere once they move to production. For example, the NAIRR should not compete with commercial vendors for the delivery of such applications. By contrast, a process should be developed in which user-created services that support AI R&D on the NAIRR can become part of the persistent infrastructure, and administered and supported alongside other existing NAIRR services, particularly if they prove valuable to multiple researchers.

Recommendation 4-15: Computing resources should be deployed using a phased approach, such that not all resources will be acquired in the same year.

The notion of mixing resources—both production versus experimental and commercial versus on premises—and acquiring them through competitions on a regular basis, will be key to keeping the NAIRR current and vibrant. A phased deployment will allow new cutting-edge resources to be deployed as existing ones age, and will "future proof" the NAIRR by ensuring that the best technologies are always part of the mix. This approach also allows the NAIRR to innovate and evolve over time.

Recommendation 4-16: The NAIRR should support federation of external edge computing resources.

The NAIRR operating model should support federation of user-supplied computing or sensors at the edge. For instance, a testbed of control towers for autonomous drones flying wireless cameras should be able to interface with the NAIRR to process data. A federated model will be more effective than trying to anticipate and build in-house, NAIRR-owned hardware for particular use cases. Projects such as SAGE⁵² might find great value in federating their own edge resources with the NAIRR.

Testbeds

AI testbeds are simulated, live, or blended environments that support research, prototyping, development, and testing of AI applications that are robust and trustworthy. The concept of a testbed can encompass the environment itself—hardware and software—as well as the data sets and frameworks that support evaluation, and the talent needed to manage the resource.

Testbeds can accelerate AI research by providing virtual or physical environments to test, simulate, explore, and develop AI. They can spur innovation in specific areas, provide opportunities to benchmark and check the quality of research, and foster cross-disciplinary collaboration.⁵³ Expanding efforts to centrally catalog AI testbeds⁵⁴ could increase accessibility to researchers and support Federal efforts to assess testbed needs. The NAIRR has an opportunity to act as a hub, cataloging and making available existing testbeds and test sets to accelerate research, increase access to inspiring testbeds, and broaden participation in AI.

Findings

Finding 4-10: Testbeds can accelerate AI research and drive advances in specific areas.

Many testbeds (including benchmarking and prototyping platforms, and living laboratories) have demonstrably driven rapid progress in areas of AI R&D, such as the MLPerf⁵⁵ benchmarking suite and test ranges created for the Defense Advanced Research Projects Agency (DARPA) Autonomous Vehicle Challenge.⁵⁶ A number of existing testbeds, such as the Georgia Tech Robotarium⁵⁷ and NIST Robotics Test Facility⁵⁸ for testing and development of physical robotics systems, POWDER-RENEW⁵⁹ and the Truman Platform⁶⁰ for development of networking and communications technologies (i.e., spectrum allocation, 5G), the NIST Facial Recognition Vendor Test⁶¹ for testing models involving such biometrics, COSMOS⁶² for assisting in developing and testing cyber-physical and smart infrastructure systems, the Energy Smart Data Center⁶³ testbed for testing models in the energy sector, CyberNET⁶⁴ and Dioptra⁶⁵ for testing cybersecurity applications, the National Transportation Research Center⁶⁶ for testing models relevant to the transportation sector, and AERPAW⁶⁷ for testing models related to unmanned aerial systems, could be beneficial to NAIRR users.

Finding 4-11: Testbeds can increase equitable involvement in AI research.

Access to testbeds can provide more researchers in less well-funded institutions with the opportunity to try new approaches for solving important problems.

Finding 4-12: Testbeds can support quality assessment of AI.

AI researchers have previously evaluated the performance of their new AI technologies on internal or proprietary data, making claims of improved performance difficult to measure or replicate. Standardized competitions can help prevent this occurrence.

Finding 4-13: Testbeds can inspire participation in AI research.

If students from kindergarteners through PhD candidates hear primarily about uninspiring, harmful, or problematic AI use cases (e.g., advertising technology, surveillance, or social media manipulation), there is a risk of deterring a diverse group of minds from considering working in or supporting the field of AI. Larger public works projects to improve safety, education, health, and equity might increase enthusiasm for working in the field. Testbeds such as those created to support the DARPA Autonomous Vehicle Challenge or various public health competitions can be helpful.

Finding 4-14: Testbeds can be broadly defined as serving the purpose of either comparison or validation.

A comparison testbed (e.g., M6,⁶⁸ Kaggle,⁶⁹ and DARPA Autonomous Vehicle Challenge⁷⁰) allows researchers to measure the effectiveness of new engineering, math, or algorithmic developments. Four specific types of comparison testbeds are listed in [Table 1](#). These can be run as explicit competitions (e.g., Kaggle) or can be permanent repositories of a challenge problem (e.g., ImageNet). Care is needed to ensure that competitions are designed to enhance diverse participation. While some competition designs can rule out participation by anyone who is not highly funded (e.g., the high cost of participating in self-driving challenges), other designs encourage citizen scientists to self-learn AI and provide individuals with notoriety leading to career opportunities.

Table 1. Examples of existing comparison testbeds

Type of Comparison Testbed	Description
1. Open Book Modeling	Test frameworks (e.g., ImageNet) in which a large repository of data is made publicly available, with a predetermined approach for training on one data set and final validation testing on a test set.
2. Closed Book Modeling	Test competitions (e.g., Kaggle, and M5/M6 time series) in which training data are made available to researchers for early training, then models' performance on a hidden test set is displayed on leaderboards during development, and then models are often finally evaluated against an additional hidden test set.

Type of Comparison Testbed	Description
3. Simulated Perceive-Decide-Act	Evaluation of different AI systems in a simulated environment. The advantages of simulation include cost and the ability to run many scenarios, giving better confidence intervals on performance. These reflect a larger end-to-end system of an AI that has multiple components including perception, statistical modeling, and automated decision components such as planners, optimizers, solvers, or controllers. Those decision components can come from academic disciplines such as reinforcement learning, operations research, control theory, game theory, and AI planning. Additional simulation testbeds are game-theoretic challenges such as Robocup. ⁷¹
4. Real-World Tests, Competitions, and Living Laboratories	These are similar to the simulated testbeds, except that evaluation is through real-world trials. An example of real-world competitions is the DARPA Autonomous Vehicle Challenge. Examples of living laboratories include the Illinois Autonomous Farm, ⁷² MLK Smart Corridor, ⁷³ Digital City Testbed Center, ⁷⁴ Virginia Smart Community Testbed, ⁷⁵ and Georgia Tech Robotarium. ⁷⁶ These can be a source of inspiration and avoid the dangers of artifacts caused by simulation. However, they are typically expensive and more difficult for conducting fine-grained comparisons of methods.

A validation testbed allows developers to decide whether an end-to-end system is acceptable to move up the maturity cycle to a more advanced phase of development.

Recommendations

Recommendation 4-17: The NAIRR should facilitate access to testbeds for AI R&D.

Testbeds are important tools for AI R&D. The NAIRR staff should include testbed experts who facilitate the development and maintenance of an up-to-date catalog of testbeds. This team would also develop and maintain a presence within the NAIRR web portal to provide access to testbed resources and advocate for consistency among testbeds under development to support reusing infrastructure and data models.

Recommendation 4-18: The NAIRR should prioritize access to comparison testbeds in order to accelerate AI research, and include infrastructure to support open-book modeling, closed-book modeling, and simulated perceive-decide-act comparison testbeds.

Comparison testbeds are quite valuable for foundational, use-inspired, and translational AI R&D, and should be prioritized within the NAIRR. On the other hand, validation testbeds are more relevant for turning AI systems into products and should not be the focus of the NAIRR. Public sector organizations, such as NIST, or industry consortium organizations are currently well positioned to sponsor validation testbeds. While real-world test, competition, and living laboratory comparison testbeds may be scarce because of their size, complexity, cost, or specialization, they should be made accessible through the NAIRR when feasible.

Recommendation 4-19: NAIRR staff should ensure that all AI comparison testbeds, including real-world tests, competitions, and living laboratories, are cataloged and made accessible to as wide a group as possible.

When Federal agencies produce requests for proposals that require testbeds, the agencies should, when possible, be aided by the NAIRR testbed staff. This aid would provide significant opportunities for consistency in setting up test mechanisms, thus avoiding duplication of effort for AI testbeds among Government agencies that fund AI research.

Testbeds that facilitate experimentation for real-world, use-inspired, and translational AI research can be highly inspiring. For example, a testbed for helping to mitigate or predict natural disasters, helping with healthcare or education, or providing test environments for field testing of mobile robots would all be visible AI use cases in which researchers and students could see the potential benefits of AI technologies and applications. Such testbeds also provide ground truth for informing research priorities: Real testbeds with real users can focus research on the most significant practical barriers to deployment. For example, an arena based on intense computer vision from thousands of cameras might be limited not by advances in computer vision theory so much as by low power cameras or network bandwidth. Additionally, simulated worlds capable of supporting thousands of researchers better support the goal of rapidly scaling AI research than testbeds using physical environments. An authoritative list of existing real-world tests, competitions, and living laboratories available to researchers would be an important contribution.

User Interface

There are many potential stakeholders, including individuals who will use the NAIRR to conduct research, data providers who will use the NAIRR to upload and manage their data sets, resource providers who will use the NAIRR portal to provide access to their infrastructure or services, and NAIRR staff who will work as managers, support personnel, and liaisons between different communities, in the AI R&D community who could become NAIRR users. The design and function of the "portal" that provides access to all NAIRR resources should be user friendly, integrate across an array of resources (including cloud and edge), and meet the needs of a varied set of stakeholders. While the portal will be one way to access NAIRR resources, it should not be the only way. Alternate access methods (e.g., shell, scripting) for more advanced users should also be available.

To provide access to a set of federated resources that will function as the NAIRR, the TF considered what attributes, design elements, and user support would be needed to integrate resources and support a dynamic research ecosystem for AI R&D. The TF also explored the feasibility of linking a heterogenous set of assets, as well as examples of organizations that provide these types of services. The findings and recommendations included here describe some attributes for the NAIRR portal, given the proposed user base, as well as general needs or design elements that should be considered in building access to an AI R&D ecosystem.

Findings

Finding 4-15: The NAIRR has the opportunity to move the broader community toward a more data-centric cyberinfrastructure ecosystem by supporting broad adoption of best practices in scientific data management, discovery, access, and curation.

A portal design should encourage community development of new capabilities. While there are many examples of technical integration across varied compute resources, integration of AI data repositories, edge computing resources, and AI testbeds is less mature and represents new elements to the cyberinfrastructure ecosystem requiring special attention. The NAIRR has the opportunity to move the broader community toward a more data-centric cyberinfrastructure ecosystem by supporting broad adoption of best practices in scientific data management, discovery, access, and curation. Edge computing applications and AI testbeds are important drivers for this transformation.

Finding 4-16: The design and function of a user portal must answer to the needs of a wide range of stakeholders.

Portal design is most effective when it is user centered, providing easy access and navigation to a wide range of users and stakeholders with reasonable responsiveness and turnaround time. Furthermore, to allow users with varied levels of experience to find, access, and use the features that are most relevant to the work they are pursuing, a portal would need to integrate seamlessly across resources (including training and support services). Finally, the portal design would have to be flexible and scalable in order to allow it to evolve as the resources within grow and change.

Finding 4-17: Community-building tools can help facilitate collaborations for both research and educational purposes.

Shared workspaces and collaboration tools can support community members as they work to develop their AI research interests and experience. Community-based platforms such as Kaggle integrate educational information, coding resources, and data sets with online competitions and discussion boards to promote community engagement and collaborative problem solving. NAIRR users might also benefit from access to collaboration and community-building tools similar in nature to those found on existing platforms (e.g., meeting rooms, chat functions).

Recommendations

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.

The portal should be a modular, extensible framework for resource and information discovery and support job execution and monitoring. The portal should provide an integrated view of all NAIRR resources but may also point to other portals and interfaces that perform specific functions well (e.g., data discovery). The technical integration of federated compute resources is quite mature, and existing resources may provide strategies and solutions for the NAIRR to emulate. Examples include the XSEDE program, the Open Science Grid (OSG),⁷⁷ and the CloudBank pilot.⁷⁸ Additionally, the portal should provide access to information and resources for a range of users, including those using the NAIRR in educational contexts, as well as the ability to submit

jobs and manage user resource allocations. Regardless of their level of experience using AI resources, NAIRR users will need to be able to select their AI application, compute resource, and data source(s) from a curated catalog and launch and monitor jobs with a minimum of effort from a portal that provides a uniform, integrated view.

Recommendation 4-21: The NAIRR should leverage user portal concepts from existing state-of-the art approaches.

The NAIRR should integrate a variety of computing, data, edge, and testbed resources while providing multiple pathways or interfaces to the resources. Access to NAIRR resources should be designed as "walk-up tooling," in which users can choose the best interface to suit their application needs and technical expertise but still have a consistent user experience across private, multi-, and hybrid-cloud infrastructures. The system should support alternate access methods (e.g., shell, scripting) for more advanced users. Finally, the user portal should be periodically subjected to a usability analysis by user interface/user experience (UI/UX) experts.

The NAIRR could outsource the design, construction, and maintenance of the user portal to a commercial entity that has previously created successful user portals. The NAIRR management entity will need to evaluate the cost effectiveness of building the user portal in house versus acquiring it commercially.

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.

The NAIRR should adopt a modular design and agile development philosophy. The NAIRR would appear to users as a set of discoverable resources supporting standard access protocols and APIs. The NAIRR should strive to balance allowing users to choose the best interface to suit their application needs and technical expertise with having resource providers adhere to standard APIs and/or best practice methodologies wherever they exist. Collectively, the federated resources should provide access to an array of computing modalities including interactive, batch, on-demand, always-on, and composable systems. The NAIRR should integrate cloud services and open source technologies whenever appropriate.

From a technical integration standpoint, the NAIRR should avoid one-off integrations, but rather encourage the emergence of standard interfaces to edge and data resources by partnering with best practice efforts. The NAIRR should develop federated data search and access services with API support, again in partnership with best-of-breed efforts.

Recommendation 4-23: Technical support and training materials appropriate for different skill levels and community-building features should be integrated into the NAIRR user portal.

Technical support for the NAIRR should be tiered, including frontline support that is built into the user interface. This in-portal help should be tightly integrated with system functions. There should be extensive "How do I use the portal?" documentation, self-help and self-paced tutorials, as well as a user forum to support community-based conversations.

The user portal should also include chat functions, meeting rooms, forums, etc., to support collaboration and community building. This strategy has been used successfully by the four regional Big Data Innovation Hubs⁷⁹ established by NSF to build a national data science community of practice with emphasis on collaborative network building and broadening participation around societal and scientific grand challenges. The user interface should be robust to the evolution of research needs, including the number of researchers, amount of data, providers of data, and pace of developing technologies.

Educational Tools and Services

The TF worked to define the scope of "education and training" in the context of the vision for the NAIRR, particularly its objective to democratize participation in AI R&D. The TF agreed that general education of AI and compute expertise is outside the scope of the mission of the NAIRR. The NAIRR should be viewed as a trustworthy platform for community building among students, researchers, resource providers, and other users. The NAIRR platform can provide facilitatory functions for educational efforts, but the NAIRR management entity should not be responsible for developing general or discipline-specific educational content.

Technical training and support related to the use of the NAIRR is within scope, and various levels of training will be required to support the array of stakeholders with access to NAIRR resources. In terms of education and training, the following findings and recommendations describe the potential training needs for NAIRR users, needs in the realm of subject matter expertise for software development and IT operations, and user support services that the NAIRR may need to provide to effectively support a broad user base.

Findings

Finding 4-18: Training for users at different skill levels (e.g., novice, intermediate, and advanced) will be needed, with room to expand the topic spaces as new resources are added to the NAIRR.

Educational resources and training options can span a range of formats, including web pages, tutorials, webinars, online training, and customized remote workshops. Hands-on or interactive training programs may be more effective for first-time users, especially for novice users and those from traditionally underserved communities. For NAIRR users, training on use of the portal itself will be needed, in addition to training or additional information on the particular resources available via the NAIRR portal.

Finding 4-19: Resource providers already create training content and tutorials.

Curating and facilitating access to existing training content generated by NAIRR resource providers is likely the easiest and most effective way to provide access to relevant, up-to-date training materials. A number of existing federated computing systems, including the NSF-funded Partnership to Advance Throughput Computing (PATH)/OSG,⁸⁰ XSEDE,⁸¹ and CloudBank,⁸² and DOE Office of Science user facilities curate training materials based on information that members or resource providers develop, making these materials available in a single location.

Finding 4-20: Access to discipline-specific technical expertise and advanced software development and IT operations (development operations) support will be a challenge.

A key challenge for the NAIRR will be to democratize access to technical expertise. A NAIRR user support center may be a mechanism for implementing this democratization, but staffing could be limited and unable to effectively support users with discipline-specific use cases. Support at this level can be viewed as a two-part problem: First, a specific type of expertise is needed to map the research problem to a technical solution, requiring discipline-specific expertise (e.g., genome analysis). Second, running the hardware and software implementations of the technical solution, making sure resources are up and running and available to users, is another, separate set of skills and expertise. User support for the NAIRR will require at least these two different kinds of expertise, both of which will be in high demand. It may also be advantageous to consider including more advanced development operations support in resource provider contracts to assist users in developing resources for specific use cases.

Recommendations

Recommendation 4-24: The NAIRR should facilitate the identification, curation, and cataloging of AI- and NAIRR-related education and training materials appropriate for different skill levels.

The NAIRR should leverage the research community and resource providers to identify and curate appropriate education and training materials for different skill levels. As indicated in the User Interface section of this chapter, front line user support should be integrated into the NAIRR by leveraging and linking to content from NAIRR resource providers. The NAIRR management entity should also build a consolidated, searchable catalog of these training materials so that everything is listed in one place. Moreover, it should facilitate identification and curation of additional AI- and resource-related training materials by the user community. The system should be instrumented to track highly used pages and tutorials in order to help the content providers better understand how users are getting the information they need and to refine how the content is delivered (e.g., static documentation versus interactive tutorials). This understanding will also help to identify best-of-breed efforts and reduce unnecessary duplication. Tiered user training documentation and interactive tutorials (beginner, intermediate, and advanced) should be created and kept current by resource providers.

Recommendation 4-25: The NAIRR should facilitate relevant educational activities but not take on new content development.

The NAIRR staff should not be responsible for developing general or discipline-specific educational content; however, the NAIRR staff should provide a platform for educational activities using NAIRR resources. For example, the NAIRR can provide a place to access educational infrastructure made available by educational resource providers. An example of this concept can be found in CloudBank, which provides users access to the Berkeley Data Stack,⁸³ a collection of tools and resources that support data science research and education at UC-Berkeley. The Berkeley Data Stack is a turnkey solution for data science education in the classroom. It provides each student with an interactive learning environment via a Jupyter notebook interface to Jupyter Books (textbook analog) developed by the instructor. It also integrates multiple containerized software components running in public clouds or on a university's compute resources (residential cloud).

Recommendation 4-26: The NAIRR management entity should provide several levels of user support including help desk and solution consulting, and incentivize community-based support.

NAIRR staff should include user portal developers, training coordinators, and data analysts reporting on NAIRR usage patterns and trends. The user support staff should include application consultants well versed in the use of public clouds for AI research (i.e., solution consultants) and education (e.g., classroom teaching). NAIRR resource providers could be connected to a network of support staff who are responsible for most user-facing functions as well as for reporting to NAIRR stakeholders and funders.

5. System Security and User Access Controls

The cybersecurity threat landscape is rapidly changing and adapting as new actors, attack methods, and vulnerabilities emerge. AI models themselves present new types of security vulnerabilities, such as adversarial examples designed to trick AI systems and model-inversion attacks that aim to extract potentially sensitive data used to train an AI model. AI research, as an asset to economic growth and national security, is a high-value target.

The threat surface for cybersecurity extends beyond technical exploits to the human element. Creating a culture of usable security⁸⁴ and training is key to mitigating human mistakes that can lead to compromise. Just as convenience could conflict with security, fostering an open research environment has tradeoffs with providing secure access to high-value information and resources and protecting intellectual property. Not only is the security vulnerability landscape rapidly changing, but so are the cybersecurity solutions and policy environment.

The NAIRR must secure the research, data, resources, and safety of its users to provide an effective and trusted research environment. The TF considered how to best enable a usable, secure, and trustworthy NAIRR; key findings and preliminary recommendations follow.

Findings

Finding 5-1: Extensive public and private sector work on security processes and policies for cyberinfrastructure can be leveraged for new systems.

Examples include the Federal Risk and Authorization Management Program (FedRAMP), NIST and International Organization for Standardization (ISO) standards, and sector-specific frameworks such as those developed by Trusted CI.⁸⁵ FedRAMP,⁸⁶ a program that promotes the adoption of secure cloud services across the Federal Government, provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMP eliminates duplicative efforts by providing a common security framework. FedRAMP is mandatory for all executive agency cloud deployments and service models, creating a large degree of familiarity and experience among government entities and current government contractors.

Finding 5-2: The constantly changing security landscape calls for an appropriately compensated staff of experts.

Cybersecurity is a rapidly changing field. New threats evolve constantly that impair defenses or introduce new risks. Security controls and mechanisms need to keep up with this rapid pace of change. Without an adept, trained, and sufficiently compensated staff who can keep up with evolving requirements, security objectives will not be met.

Finding 5-3: Security risks to cyberinfrastructure are as much human as technical.

In addition to technical controls, any cyberinfrastructure requires clear, transparent policies about permitted system users, accesses, and technologies; user training on security risks and policy compliance; and mechanisms for monitoring and enforcing security policies and updating them as needed.

Recommendations

Recommendation 5-1: Use FedRAMP standards (not FedRAMP approval) during a transition period of 2 years, while developing a NAIRR-specific approval process consistent with applicable laws and regulations (e.g., when the NAIRR facilitates access to government data).

The NAIRR management entity should assemble an advisory group to recommend security standards for all resources provided through the NAIRR. As part of its work, in the case of NAIRR cloud resources, the advisory group should develop a NAIRR-dedicated FedRAMP-inspired solution and plan for regular monitoring and updating. The advisory group should have expertise in security, operations, risk management, and scientific research considerations (e.g., security versus usability). Security policies should focus on incentive structure rather than on enforcement structure. Regular monitoring and update of the security controls should be incorporated, including considerations of red-teaming system security protocols, as part of NAIRR governance.

Recommendation 5-2: The NAIRR should implement system safeguards using zero trust architecture as well as the Five Safes framework: safe projects, safe people, safe settings, safe data, and safe exports.

Zero trust architecture presumes that no actor, system, network, or service operating outside or within the security perimeter is trusted.⁸⁷ U.S. Federal agencies are moving toward zero trust architecture by the end of fiscal year 2024. This strategy emphasizes the importance of strong identity and access controls, including multi-factor authentication, as well as defenses against phishing.

The Five Safes framework⁸⁸ was originally developed in the United Kingdom and is now used internationally. The framework structures protection across five dimensions: research projects and individuals working on projects are reviewed and approved; people using the resource must sign security agreements and complete training, and users' access is monitored; settings operationalize security needs and are managed through a central platform; data is appropriately safeguarded against security, reidentification, and privacy risks; and exports are technically and contractually controlled, and evaluated and monitored to prevent unauthorized disclosure.⁸⁹ For data not already covered by existing Federal regulatory frameworks, the NAIRR will need to make substantial investments in establishing technical controls as well as governance to determine the policies for how data are used within the NAIRR.

Recommendation 5-3: If the NAIRR includes confidential or sensitive data, the NAIRR management entity should implement a tiered access model to accommodate heterogeneous security needs.

Should the NAIRR provide access to confidential or sensitive data, that data will need to be designed to secure heterogeneous assets within a rapidly evolving technological and threat landscape and for a diverse user base. While some information can be open, several types of information assets (e.g., personally identifiable information (PII) or other sensitive data) require higher levels of protection. For example, data collected on human subjects is often confidential. Even if PII is not directly included, there is still a risk of the data being used to reidentify individuals and cause harm.⁹⁰ Access to statutorily protected government data, from either statistical agencies or programmatic sources, will be subject to Federal regulatory and statutory frameworks and standard access procedures; confidential business or intellectual property-protected data will also have security needs. Additionally, the compute environment itself will have to be secured to avoid unauthorized use.

The heterogeneity in protection requirements will require different levels of permissions for individuals to work at different levels of access. All users should access NAIRR resource components that are generally available through a single sign-on, with higher levels of access available to users performing sensitive research. This approach is consistent with that recommended by the Commission on Evidence-Based Policymaking for managing risks associated with the sharing of government data.

The NAIRR should provide multiple levels of access, from a "green-zone" open to all users to secure enclaves for users working with sensitive data, if included, or algorithms. For cloud resources within the NAIRR, FedRAMP standards should be viewed as the minimum acceptable security requirements, with additional protections encouraged. Full FedRAMP approval may be required for certain information assets as a matter of law or regulation—namely, those that interact with sensitive government-held data. Resources with lower risk levels could adopt just the needed aspects of FedRAMP standards. The high bar of FedRAMP approval could be prohibitive to potential resource providers. A flexible, tiered approach can help open participation to different resource providers while supporting high security where needed.

Recommendation 5-4: Provide regular and continuous hands-on training for NAIRR staff and users.

To reduce human error, all users should receive, at a minimum, basic, hands-on training on system security policies and processes. This training should include the following:

- a. System protections: how privacy, data, and assets will be protected in terms understandable to security non-experts. System protections training should ideally include all system and resource elements, including those managed by other entities, but at a minimum should include system protections for the elements managed by the NAIRR. This description is meant to promote user trust in the system.

- b. The user's role in the security and privacy of the system and its data: why that role is needed (potential security/privacy risks) and why that role is important (potential consequences).
- c. Step-by-step instructions that guide users through any security or privacy actions they are expected to take (e.g., if they are permitted to share or set access control on data, how they would go about doing that, and considerations for how much access should be granted and to whom).
- d. Different tiers of training appropriate to the different levels of access. Consider including exceptions for training requirements for very simple, classroom exercise-type uses in which the NAIRR is used as part of a lecture to demonstrate a concept or project.

The initial training could be offered via several different media (e.g., videos, step-through computer-based training, PDF information sheets) and be accessible to those with differing abilities. Delivery and verification of security training completion should be coordinated with delivery and verification of the privacy, civil rights, and civil liberties training described in [Chapter 6](#). Users should also have access to security-targeted help from a live user support team, and access to help at any time during their use of the system through general FAQs or other static resources describing the system security protections, policies, and processes.

Recommendation 5-5: Regular system monitoring and update and associated security controls should be included as part of NAIRR management, governance, and administration activities.

The NAIRR should embrace a "living" security plan that evolves with needs and the threat landscape, and that is continually refreshed based on emerging research results. While this is true of all aspects of the NAIRR, it is especially true of cybersecurity. Resource providers included in the NAIRR cyberinfrastructure should provide appropriate security controls for the resources they provide.

Recommendation 5-6: To accommodate evolving needs, the NAIRR should have dedicated technical security staff experts who can keep up with evolving requirements, are paid sufficiently, and are employed directly by the NAIRR management entity.

Ensuring sufficient pay to attract needed technical expertise is a crosscutting issue, applicable to supporting many of the functions of the NAIRR. The solution ultimately may have implications for the ownership, administration, and governance of the NAIRR.

6. Privacy, Civil Rights, and Civil Liberties

Recent high-profile examples of bias in AI⁹¹ and the use of AI applications in ways that exacerbate inequality⁹² or impact human rights⁹³ have raised public awareness of how personal data is being collected, and how AI is being used to inform and personalize key aspects of everyday life including hiring, financial services, and digital communications. Further, responses submitted as part of the RFI⁹⁴ on the NAIRR also indicate that key areas of interest associated with the use of AI include privacy and civil liberties, safety, and oversight of AI R&D. The NAIRR must carefully address appropriate safeguards for privacy, civil rights, and civil liberties, while at the same time fostering intellectual curiosity that can pay dividends for our Nation's long-term competitiveness.

To earn and maintain public trust, research in areas that may impact privacy, civil rights, or civil liberties will need to be reviewed, approved, and performed in a way that meets the expectations of civil society and protects subjects' rights. The public and civil society groups will likely want to understand where and when AI developed using the NAIRR is being used, as well as obtain assurances that models developed by NAIRR users are guiding decisions that are both equitable and fair. A study currently underway at the National Academies of Sciences, Engineering, and Medicine will help to inform future discussions on this topic.⁹⁵

The TF leveraged the RFI responses, panel discussions from TF meetings, and additional conversations with experts to understand issues, requirements, and potential solutions related to ensuring privacy, civil rights, and civil liberties in the design, implementation, and operation of the NAIRR. The findings and recommendations that follow reflect key themes and principles, and describe current aspects of the AI R&D landscape that the NAIRR can improve upon to help provide a research environment that integrates privacy, civil rights, and civil liberties considerations across all points in the AI research continuum.

Findings

Finding 6-1: Engaging diverse stakeholders is critical to ensure that privacy, civil rights, and civil liberties are being appropriately considered and protected.

Explanations of data sets and AI applications for non-technical readers written in plain language accessible to all communities,⁹⁶ with different reading levels, abilities, and access requirements, are needed to ensure that AI research is accessible to all the people it will impact. Further, it is essential that any review group or advisory board be broadly diverse and include expertise and representation that fully spans and reflects the diversity of the stakeholder communities anticipated for the NAIRR. Opportunities to engage communities or stakeholder groups can help ensure that issues pertaining to privacy, civil rights, and civil liberties are being carefully considered as AI research projects and applications are being designed.

Finding 6-2: The resources and tools currently used to perform AI R&D are too often designed without privacy, civil rights, and civil liberties considerations in mind.

It will be important to provide an environment that researchers who focus on AI trustworthiness can easily examine and build on the results of researchers who focus on specific applications of AI or machine learning. Ensuring fairness and preventing bias may rely heavily on interdisciplinary or collaborative research that has not been done before, and efforts to document the status, structure, and provenance of existing data sets and models, as well as mechanisms to deprecate data sets found to be problematic, may be needed.

Finding 6-3: To safeguard privacy, civil rights, and civil liberties, AI researchers need to be trained to recognize potential challenges and issues.

New applications and uses of AI are likely to have specific issues of trustworthiness and fairness that researchers cannot reasonably foresee. While it may not be feasible to require all researchers to receive in-depth training on different approaches to mitigate bias and promote trustworthiness, many existing training resources include information on fairness, ethics, and social responsibilities associated with AI research. This information may be employed to help researchers understand their obligations in performing responsible and ethical AI research.

Finding 6-4: To effectively identify, monitor, and address issues that may challenge privacy, civil rights, or civil liberties, oversight processes must be integrated into the administration and operation of AI R&D programs.

Ensuring that AI R&D is ethical requires strong governance frameworks and stewardship across the full complement of capabilities which drive that work, including computing, data, and software. In particular, it is important that governance structures and associated operational processes fundamentally address privacy, civil rights, and civil liberties principles in all aspects of AI R&D.

Recommendations

The NAIRR has an opportunity to serve as an exemplar for how transparent and ethical AI R&D can be performed with proper training and oversight at multiple levels. The recommendations in this section reflect TF thinking on how the NAIRR can support AI R&D while keeping privacy, civil rights, and civil liberties considerations at the forefront.

Recommendation 6-1: Transparency, access for diverse users, and oversight should drive the efforts of the NAIRR to protect privacy, civil rights, and civil liberties.

Efforts to ensure that NAIRR operations, research, and governance are conducted in a transparent fashion with appropriate oversight should be integrated across all aspects of the design, implementation, administration, management, and use of the NAIRR. There should be a clear vetting and evaluation process to receive approval for any research to be done on the NAIRR. In addition, policies and requirements to protect privacy, civil rights, and civil liberties should be integrated into the user agreements governing access to all NAIRR resources. Operational protocols for the NAIRR should include mechanisms to support security and privacy requirements, enable monitoring, and ensure compliance with all governance policies. In addition, all

stakeholders should be subject to regular reporting requirements, and these reports should be made publicly available. Finally, the NAIRR governance should have established protocols, a communication infrastructure, and an engagement plan to follow if a violation is identified, with an emphasis on transparently communicating to the public what happened, any corrective actions taken, and efforts to prevent future issues.

Recommendation 6-2: The NAIRR should establish an ethics review process to vet all resources included in the system and the research performed within.

External ethics reviewers (as generally described in [Chapter 3](#) of this report) should be leveraged for this purpose. While the majority of data in the NAIRR is not expected to have ethical concerns, the NAIRR management entity should establish and implement acceptance criteria and recommended best practices for all resources joining the NAIRR to ensure that they are vetted from privacy, civil rights, civil liberties, and inclusivity perspectives. This acceptance criteria should be more stringent for resources that are likely to be used in contexts that raise heightened concerns about privacy, civil rights, and civil liberties. In the case of third-party data sets made available via the NAIRR, this vetting process would need to be developed and could include establishing certification standards and/or providing trusted and validated reference data sets for testing (i.e., as an audit system). Only after appropriate vetting may these data sets be included in the NAIRR. In addition, the inclusion of higher risk data sets that have been modified with embedded privacy protections must be reviewed by potentially affected communities, because of the possible impact on those communities.

In the case of research supported by Federal funding agencies, existing mechanisms such as institutional review boards may be leveraged to ensure proper review and vetting. For research that is supported through other means, the funder or the NAIRR management entity should integrate guidance pertaining to privacy, civil rights, and civil liberties into the evaluation process. The outcomes from research enabled by the NAIRR should also be vetted at regular intervals and over the long term, to the extent feasible, to ensure that the research is not leading to violations of privacy, civil rights, and civil liberties.

Recommendation 6-3: NAIRR users should be required to complete training modules before being granted access to the NAIRR, and this training should be refreshed annually.

Ensuring awareness about rights, responsibilities, and best practices related to privacy, civil rights, and civil liberties is essential. As a result, all NAIRR users should be required to take training before being granted access to the NAIRR, and this training should be refreshed annually. The level of training required should be commensurate with the nature of NAIRR usage. For example, short-term, non-sensitive use of the NAIRR, such as a short classroom exercise, may warrant less rigorous requirements. The training may be delivered online, and it should be developed in conjunction with experts in equitable and trustworthy AI with specialization in sensitivity training, and in civil rights and civil liberties and their application to AI research. Because the user base for the NAIRR is expected to be broad and diverse, training should be tailored for various audiences.

Recommendation 6-4: NAIRR resources should be allocated to specifically support research on AI trustworthiness and should develop best practices for working with flawed models and information.

Ensuring transparency regarding AI data sets, common practices, and decisions that inform development of use cases is essential, and, in this regard, the NAIRR can serve as a leader in establishing and promulgating best practices. The NAIRR should maintain an inventory of provided data sets and their history and provenance using Persistent Digital Identifiers, including a list of deprecated data sets. The NAIRR should also foster the curation of biased-by-design data sets to help advance research on AI bias, and to test models for robustness. The NAIRR should also enable researchers who focus on fairness to examine data sets and applications, and build upon the results of research enabled by the NAIRR. By supporting a community of practice for researchers focused on fairness and supporting development of auditing tools and mechanisms, the NAIRR can support efforts to incorporate privacy, civil rights, and civil liberties considerations into the design of new AI R&D resources and tools.

7. Next Steps for the Task Force

The preliminary findings and recommendations presented in this interim report outline the NAIRR TF's current vision for the NAIRR. The objective for the interim report is to develop this vision for the NAIRR and its key components, *not* to go into details of how the envisioned NAIRR would be implemented. The interim report will be followed by a final report in December 2022 that will provide a more detailed roadmap and implementation plan for realizing this vision.

Moving forward, the TF will refine its findings and recommendations for the design of the NAIRR and deliberate on remaining open questions. In doing so, the TF will develop recommendations related to the specific steps needed to build the NAIRR and realize the TF's strategic vision for the NAIRR: *to strengthen and democratize the U.S. AI innovation ecosystem in a way that protects privacy, civil rights, and civil liberties.*

At this time, the TF welcomes public feedback on the NAIRR features proposed herein or remaining open questions. To share your input, please reply to the Request for Information published in [Federal Register Notice 2022-11223](#) by June 30, 2022, or participate in the public listening session on June 23, 2022, by following the instructions provided in [Federal Register Notice 2022-11222](#). In the months ahead, the TF will integrate this public feedback with other inputs (for example, through additional panel discussions at upcoming TF meetings) and further deliberate to arrive at a final report outlining a detailed roadmap and implementation plan for the NAIRR. The TF's final report is anticipated to be released in December 2022.

Appendix A. Definitions

AI Testbeds: Simulated, live, or blended environments that support prototyping, development, and testing of AI applications that are robust and trustworthy. The concept of a testbed can encompass the environment itself—hardware and software—as well as the data sets and frameworks that support evaluation, and the talent needed to manage the resource.

Artificial intelligence (AI): See [Box 1](#).

Comparison Testbed: A category of testbeds that allows researchers to measure the effectiveness of new engineering, math, or algorithmic developments. These can be run as explicit competitions (e.g., Kaggle) or exist as permanent repositories of a challenge problem (e.g., ImageNet).

Cyberinfrastructure: Infrastructure based on distributed computer, information, and communication technologies, including the enabling hardware, algorithms, software, services, communications, institutions, and expertise.

Experimental System/Resource: A system or resource that is exploring a new hardware or software capability and may provide an immature or rapidly evolving environment for the user to run in. Users may expect additional efforts to port applications to properly use the capabilities of the system, rather than a "turnkey" environment, and not all use cases will be well supported.

Federated system: A set of semi-autonomous, decentralized resources that agree to a standard set of protocols allowing for integration, interoperability, and information sharing.

Foundational AI research: Discovery-oriented fundamental research, including knowledge representation, reasoning, planning, learning, language processing, perception, vision, and motion and manipulation, that seeks to advance the frontiers of AI.

Fundamental Research: Also known as basic research, research that spans the full spectrum from foundational, discovery-oriented to use-inspired, solution-oriented research.

National AI Research Resource (NAIRR): See [Box 2](#).

Production system/resource: A computing system or resource on which users can run state-of-the-art software tools, frameworks, and applications without modification and with reasonable expectations of stability, reliability, and quality of service. A system in production should be expected to have passed a set of predefined acceptance tests that measure performance, usability, and stability of the environment.

Research on AI: Foundational, use-inspired, and translational research that advances scientific understanding of the nature of intelligence, mathematical understanding of the behavior of adaptive/autonomous systems, or algorithmic understanding of techniques in the component areas of AI (including perception, learning, planning, and robotics), as well as research related to robustness, scalability, reliability, safety, security, privacy, interpretability, and equity of AI systems.

Testbeds: Platforms used to conduct research and validate theories, tools, or technologies in a rigorous, replicable manner.

Translational AI research: Research that bridges foundational and use-inspired research with the delivery and deployment of its outcomes to the target community, and supports essential bi-direction interplays by which delivery and deployment process informs the research; e.g., translating research results from the lab to the market and society.

Use-inspired AI research: Fundamental research in AI that is motivated or inspired by particular use cases and seeks to advance both the frontiers of AI and specific use cases.

Validation testbed: A category of testbed that allows developers to decide whether an end-to-end system is acceptable to move up the maturity cycle to a more advanced phase of development.

Appendix B. Briefers to the Task Force

The Task Force (TF) held seven public meetings between its launch in July 2021 and the release of this interim report. At these meetings, the TF discussed and developed a vision for the NAIRR, heard input from invited expert speakers and panelists, and deliberated on key findings and preliminary recommendations for the design of the NAIRR. These outside expert briefers and panelists, along with their affiliations, are listed here.

July 28, 2021

The STRIDES program

Andrea Norris and Nick Weber, *National Institutes of Health*

August 20, 2021

Value proposition and intended outcomes of a NAIRR

Damian Clarke, *Chief Information Officer and Computer Science Faculty, Alabama A&M University (now Chief Information Officer, Meharry Medical College)*

James Deaton, *Executive Director, Great Plains Network*

Deborah Dent, *Chief Information Officer, Jackson State University*

Tripti Sinha, *Assistant Vice President and Chief Technology Officer, University of Maryland and Executive Director of the Mid-Atlantic Crossroads (MAX)*

Talitha Washington, *Director, Atlanta University Center Consortium Data Science Initiative*

Ownership, governance, and administration models

Sharon Broude Geva, *Director for Innovation and Computational Research, University of Michigan*

Manish Parashar, *Office Director, Office of Advanced Cyberinfrastructure, National Science Foundation⁹⁷*

Gina Tourassi, *Director, National Center of Computational Sciences and the Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory*

John Towns, *Executive Associate Director for Engagement, National Center for Supercomputing Applications and Deputy CIO for Research IT, University of Illinois at Urbana-Champaign*

Frank Würthwein, *Executive Director, San Diego Supercomputer Center*

October 25, 2021

Data resources

Ian Foster, *Director, Data Science and Learning Division, Argonne National Laboratory; Professor of Computer Science, University of Chicago*

Robert L. Grossman, *Professor of Medicine and Computer Science, University of Chicago*

Ron Hutchins, *Vice Provost for Academic Technologies, University of Virginia*

Anita Nikolich, *Research Scientist and Director of Research and Technology Innovation, University of Illinois at Urbana-Champaign*

Nancy Potok, *CEO, NAPx Consulting; former Chief Statistician of the United States*

Andrew Trask, *Leader, OpenMined*

User resources: portal interface, educational tools

Tiziana Ferrari, *Director, EGI Foundation*

Kimberly Greene Starks, *Global Lead, Infrastructure and Technology Strategy, IBM University Programs*

Ana Hunsinger, *Vice President for Community Engagement, Internet2*

Ed Lazowska, *Professor and Bill & Melinda Gates Chair Emeritus, Paul G. Allen School of Computer Science & Engineering University of Washington*

December 13, 2021

Privacy, civil rights, and civil liberties requirements

Solon Barocas, *Principal Researcher, Microsoft Research; Adjunct Assistant Professor, Information Science, Cornell University*

Lujo Bauer, *Professor, Electrical & Computer Engineering and Computer Science, Carnegie Mellon University*

danah boyd, *Partner Researcher, Microsoft Research; and Founder/President, Data & Society*

Deborah Raji, *Fellow, Mozilla Foundation*

Nicol Turner Lee, *Senior Fellow and Director of the Center for Technology Innovation, Brookings Institution*

Hannah Quay-de la Vallee, *Senior Technologist, Center for Democracy and Technology*

February 16, 2022

User perspectives on the NAIRR

Tom Dietterich, *Distinguished Professor Emeritus in Computer Science, Oregon State University*

Susanta Ghosh, *Assistant Professor in Mechanical Engineering-Engineering Mechanics, Michigan Technological University*

Kinnis Gosha, *Hortinius I. Chenault Endowed Associate Professor of Computer Science, Morehouse College*

Gail Rosen, *Professor, Drexel University*

Rima Seilova-Olson, *Co-Founder and Chief AI Scientist, Kintsugi*

Carlos Theran, *Research Associate, Florida A&M University*

April 8, 2022

Building responsible AI review processes for the NAIRR

Beena Ammanath, *Author, Trustworthy AI and Head of Global Deloitte AI Institute*

Michael Bernstein, *Associate Professor of Computer Science, Stanford University*

Arvind Narayanan, *Associate Professor of Computer Science, Princeton University*

Beth Plale, *Michael A. and Laurie Burns McRobbie Bicentennial Professor of Computer Engineering and Executive Director, Pervasive Technology Institute, Indiana University Bloomington*

Christo Wilson, *Associate Professor of Computer Science, Northeastern University*

May 20, 2022

No external speakers; the purpose of the meeting was to vote on the interim report.

Appendix C. Public Input Provided in Response to the Federal Request for Information

A Request for Information on the design of a NAIRR was posted in the Federal Register (86 FR 39081) on July 23, 2021; the comment period closed on October 1, 2021. The TF received 84 responses. The list of respondents to this Request for Information follows; the full text of the responses is available at <https://www.ai.gov/nairrtf/86-fr-39081-responses/>.

- Abas Abdoli, Ryan N Coffee, Auralee Edelen, Michael Kagan, Daniel Ratner, Sohail Reddy, and Kazuhiro Terao
- Accenture
- ACM U.S. Technology Policy Committee
- AI Now Institute of New York University and Data & Society Research Institute
- AI Redefined, Inc.
- Aishik Ghosh
- Amazon Web Services
- American Civil Liberties Union (ACLU)
- American Psychological Association (APA)
- Amitha Domalpally, Roomasa Channa, Pila Ossorio
- Anthropic
- Argonne National Laboratory
- Atlantic Council GeoTech Center
- BeeHero
- Ben Freed and Howie Choset
- Booz Allen Hamilton
- Cadence
- CalypsoAI Corp.
- Carnegie Mellon University
- Center for Data Innovation
- Center for Democracy and Technology
- Center for Security and Emerging Technology
- Cerner Corporation
- Computing Community Consortium, Computing Research Association-Industry, Association for the Advancement of Artificial Intelligence
- Consumer Reports
- CrowdAI
- Deloitte
- Digital Diagnostics
- Electronic Privacy Information Center (EPIC)
- Engine
- FABRIC Testbed
- Google
- Hewlett Packard Enterprise
- Hyperion Research
- IBM
- Institute of Electrical and Electronics Engineers (IEEE) Standards Association
- Indiana University
- Infiltron
- Information Technology Industry Council
- Internet2
- Jared Freeman, Drew Leins, Niall Gaffney
- John T. Feddema, David J. Stracuzzi, James R. Stewart

- Kermit Kubitz
- Lawrence Berkeley National Lab
- Lawrence Berkeley National Laboratory Machine Learning Group
- Lawrence Livermore National Laboratory
- Maria Patterson
- Mathematica
- Medical Imaging and Resource Center, University of Chicago
- Michael August
- Microsoft
- Moffitt Cancer Center
- NASA
- National Center for Atmospheric Research
- National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign
- National Energy Technology Laboratory
- NiyamIT, Inc.
- Noblis
- Northeastern University
- NSF AI Institute for Artificial Intelligence and Fundamental Interactions
- NSF AI Institutes
- NVIDIA
- Open Commons Consortium at the Center for Computational Science Research, Inc.
- Oracle America, Inc.
- Palantir Technologies, Inc.
- Partnership on AI
- Representatives from the National Oceanic and Atmospheric Administration (NOAA) Artificial Intelligence Executive Committee (NAIEC) and the Center for Artificial Intelligence (NCAI)
- S. Joseph Sirintrapun
- SAS
- Savash Kapoor, Mihir Kshirsagar, Arvind Narayanan
- Sean Ekins
- Stanford Libraries
- Stanford University Institute for Human-Centered Artificial Intelligence (HAI)
- Steve Xiao
- The Aerospace Corporation
- The Alexandria Archive Institute (Open Context)
- The Data Foundation
- The Enterprise Neurosystem
- The MITRE Corporation
- Thomas Yankeelov
- U.S. Chamber of Commerce Technology Engagement Center
- University of Florida
- University of Illinois Chicago
- Wayne Gilmore, John Goodhue, Christopher N. Hill, David Kaelli, Eric Kolaczyk, Jim Kurose, Scott Yackel

Appendix D. Subject Matter Experts Consulted by TF Members

Pete Beckman,
Argonne National Laboratory

Suzette Kent,
Kent Advisory Services

Jim Brase,
*COVID-19 HPC Consortium
and Lawrence Livermore National Lab*

Tony LaVoie,
*National Oceanic
and Atmospheric Administration*

Kate Crawford,
AINow (NYU)

Aaminah Norris,
Algorithmic Justice League

Ian Ferreira,
Core Scientific, Inc.

Joris Poort,
ReScale, Inc.

Brett Goldstein

Nancy Potok,
NAPx Consulting

Julie Haney, *National Institute
of Standards and Technology*

Adam Schwartz,
Ames Laboratory

Nick Hart,
Data Foundation

Brock Webb,
Bureau of the Census

Bob Jackson,
Spherecom Enterprises

Harlan Yu,
Upturn

Appendix E. NAIRR Task Force Staff and Contributors

Tess deBlanc-Knowles,
*White House Office of Science and
Technology Policy and National Science
Foundation*

Emily Grumbling,
*IDA Science and Technology Policy
Institute*

Matthew Ishimaru,
*IDA Science and Technology Policy
Institute*

Morgan Livingston,
*IDA Science and Technology Policy
Institute*

Lisa Van Pay,
*IDA Science and Technology Policy
Institute*

Taylor White,
*IDA Science and Technology Policy
Institute*

Appendix F.Abbreviations

ACLU	American Civil Liberties Union
AI	Artificial Intelligence
API	Application Programming Interface
DARPA	Defense Advanced Research Projects Agency
DEIA	Diversity, Equity, Inclusion, and Accessibility
DHS	Department of Homeland Security
DOE	Department of Energy
DUA	Data Use Agreement
EPIC	Electronic Privacy Information Center
FAIR	Findability, Accessibility, Interoperability, and Reusability
FedRAMP	Federal Risk and Authorization Management Program
FFRDC	Federally Funded Research and Development Center
Gbps	Gigabits per Second
HAI	Human-Centered Artificial Intelligence
HPC	High-Performance Computing
HUD	Department of Housing and Urban Development
IEEE	Institute of Electrical and Electronics Engineers
ISO	International Organization for Standardization
NAIEC	National Oceanic and Atmospheric Administration (NOAA) Artificial Intelligence Executive Committee
NAIRR	National Artificial Intelligence Research Resource
NASA	National Aeronautics and Space Administration
NIH	National Institutes of Health
NIST	National Institute of Standards and Technology
NOAA	National Oceanic and Atmospheric Administration
NSF	National Science Foundation
OMB	Office of Management and Budget
OSG	Open Science Grid
OSTP	Office of Science and Technology Policy
PATh	Partnership to Advance Throughput Computing
PII	Personally Identifiable Information
R&D	Research and Development
RFI	Request for Information
SBIR	Small Business Innovation Research
STRIDES	(NIH) Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability
STTR	Small Business Technology Transfer
TB	Terabyte
TF	Task Force
UI/UX	User Interface/User Experience
XSEDE	eXtreme Science and Engineering Discovery Environment

Appendix G. Notes

- ¹ "AI" also refers broadly to the associated field of inquiry.
- ² Rishi Bommasani et al., "On the opportunities and risks of foundation models," *CoRR* abs/2108.07258 (2021), <https://arxiv.org/abs/2108.07258>.
- ³ AlphaFold Protein Structure Database (DeepMind and EMBL's European Bioinformatics Institute), accessed April 26, 2022, <https://alphafold.ebi.ac.uk/>.
Jonas Degraeve et al., "Magnetic control of Tokamak plasmas through deep reinforcement learning," *Nature* 602, no. 7897 (2022): pp. 414-419, <https://doi.org/10.1038/s41586-021-04301-9>.
Kshitij Bansal et al., "HOList: An environment for machine learning of higher-order theorem proving," *arXiv*, 2019, <https://doi.org/10.48550/ARXIV.1904.03241>.
- ⁴ Daniel E. Ho et al., "Building a national AI research resource," (2021), https://hai.stanford.edu/sites/default/files/2022-01/HAI_NRCR_v17.pdf.
- ⁵ Steve Lohr, "At tech's leading edge, worry about a concentration of power," *New York Times*, September 26, 2019, <https://www.nytimes.com/2019/09/26/technology/ai-computer-expense.html>.
Eric Schmidt et al., "Final report," *National Security Commission on Artificial Intelligence*, 2021, <https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf>.
Nur Ahmed and Muntasir Wahed, "The de-democratization of AI: Deep learning and the compute divide in artificial intelligence research," *arXiv preprint arXiv:2010.15581* (2020), <https://arxiv.org/abs/2010.15581>.
Ryan Khurana, "How to make artificial intelligence more democratic," *Scientific American*, January 2, 2021, <https://www.scientificamerican.com/article/how-to-make-artificial-intelligence-more-democratic/>.
Nathan Benaich and Ian Hogarth, "State of AI report 2021," *State of AI*, October 12, 2021, <https://www.stateof.ai/>.
- ⁶ Schmidt et al., "Final report."
Ahmed and Wahed, "The de-democratization of AI: Deep learning and the compute divide in artificial intelligence research."
Khurana, "How to make artificial intelligence more democratic."
Benaich and Hogarth, "State of AI report 2021."
- ⁷ Benaich and Hogarth, "State of AI report 2021."
- ⁸ Kate Crawford et al., "The AI now report," (A summary of the AI Now public symposium, hosted by the White House and New York University's Information Law Institute, July 7, 2016), https://ainowinstitute.org/AI_Now_2016_Report.pdf.
- ⁹ Nine percent of U.S. resident AI PhDs graduating in 2020 were of unknown race/ethnicity.
Daniel Zhang, Nestor Maslej, Erik Brynjolfsson, John Etchemendy, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Michael Sellitto, Ellie Sakhaee, Yoav Shoham, Jack Clark, and Raymond Perrault, "The AI Index 2022 Annual Report," AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022., <https://aiindex.stanford.edu/report/>.
Stuart Zweben and Betsy Bizot, "2020 Taulbee survey," *Computing Research Association*, 33, no. 5, (2021), <https://cra.org/crn/2021/05/2020-taulbee-survey/>.
- ¹⁰ Zhang et al., "The AI index 2022 annual report."
Zweben and Bizot, "2020 Taulbee survey."
- ¹¹ Meredith Whittaker et al., "Disability, bias, and AI," *AI Now Institute* (2019), <https://ainowinstitute.org/disabilitybiasai-2019.pdf>.

-
- Joy Buolamwini and Timnit Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," In Conference on fairness, accountability and transparency, pp. 77–91, PMLR, 2018, <http://proceedings.mlr.press/v81/buolamwini18a.html>.
- Kate Crawford and Trevor Paglen, "Excavating AI: The politics of images in machine learning training sets," *Liverpool Biennial*, 9, (2019), <https://www.biennial.com/journal/issue-9/excavating-ai-the-politics-of-images-in-machine-learning-training-sets>.
- Su Lin Blodgett et al., "Language (technology) is power: A critical survey of 'bias' in NLP," *arXiv preprint arXiv:2005.14050* (2020), <https://arxiv.org/abs/2005.14050>.
- Virginia Eubanks. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press, 2018.
- Kate Crawford. "The atlas of AI." In *The Atlas of AI*. Yale University Press, 2021.
- Ruha Benjamin. "Race after technology: Abolitionist tools for the new Jim code." *Social forces* (2019).
- Safiya Umoja Noble. "Algorithms of oppression." In *Algorithms of Oppression*. New York University Press, 2018.
- Catherine D'ignazio and Lauren F. Klein. *Data feminism*. MIT press, 2020.
- ¹² Zhang et al., "The AI index 2022 annual report."
Zweben and Bizot, "2020 Taulbee survey."
- ¹³ Nathan Baker et al., *Workshop report on basic research needs for scientific machine learning: Core technologies for artificial intelligence*, U.S. Department of Energy Office of Science, Washington, DC, 2019, <https://doi.org/10.2172/1478744>.
The national artificial intelligence research and development strategic plan: 2019 update, National Science and Technology Council, 2019, <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf>.
Schmidt et al., "Final report."
Michael L. Littman et al., *Gathering strength, gathering storms: The one hundred year study on artificial intelligence (AI100) 2021 study panel report* (Stanford, CA: Stanford University, 2021), <http://ai100.stanford.edu/2021-report>.
Yolanda Gil and Bart Selman, *A 20-year community roadmap for artificial intelligence research in the US* (Computing Community Consortium and Association for the Advancement of Artificial Intelligence, August 6, 2019), arXiv:1908.02624, <https://cra.org/ccc/wp-content/uploads/sites/2/2019/08/Community-Roadmap-for-AI-Research.pdf>.
- ¹⁴ 15 U.S.C. § 9411, <https://uscode.house.gov/view.xhtml?req=granuleid%3AUSC-prelim-title15-chapter119-subchapter1&saved=%7CZ3JhbnVsZWlkOjVlcmVsaW0tdGl0bGUxNS1zZWNoaW9uOTQxNQ%3D%3D%7C%7C%7C0%7Cfalse%7Cprelim&edition=prelim>.
- ¹⁵ *The national artificial intelligence research and development strategic plan: 2019 update*.
- ¹⁶ *Recommendations for leveraging cloud computing resources for federally funded artificial intelligence research and development*, National Science and Technology Council, 2020, <https://www.nitrd.gov/pubs/Recommendations-Cloud-AI-RD-Nov2020.pdf>.
- ¹⁷ 15 U.S.C. § 9415, <https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title15-section9415&num=0&edition=prelim>.
- ¹⁸ For example, see the list of technology areas supported by the National Science Foundation's SBIR/STTR Program, America's Seed Fund: <https://seedfund.nsf.gov/portfolio/>, or at the Department of Energy's SBIR program: <https://www.energy.gov/science/sbir/small-business-innovation-research-and-small-business-technology-transfer>.
- ¹⁹ For example, the COVID-19 High Performance Computing Consortium: <https://covid19-hpc-consortium.org/>.
- ²⁰ Federal Acquisition Regulation 35.017-2: "Establishing or changing an FFRDC."
- ²¹ Federal Acquisition Regulation 35.017: "Federally funded research and development centers."

-
- ²² For example, the Dimensions platform (dimensions.ai) shows that there are over 20,000 active AI researchers in the United States based on a customized query of AI-related keywords. This number has grown significantly in the last five years and that growth is expected to continue.
- ²³ See, for example, National Science Foundation, "About public access," accessed April 27, 2022, <https://www.research.gov/research-web/content/aboutpublicaccess>.
- ²⁴ See also W.K. Kellogg Foundation, *Logic model development guide*, updated January 2004, https://www.naccho.org/uploads/downloadable-resources/Programs/Public-Health-Infrastructure/KelloggLogicModelGuide_161122_162808.pdf.
- ²⁵ For example, Zweben and Bizot, "2020 Taulbee survey."
- ²⁶ Mary L. Gray and Siddharth Suri, *Ghost work: How to stop Silicon Valley from building a new global underclass* (Eamon Dolan Books, 2019).
- ²⁷ U.S. Congress, House, *Foundations for Evidence-Based Policymaking Act of 2018*, HR 4174, Public Law 115-435, 115th Congress, <https://www.congress.gov/bill/115th-congress/house-bill/4174>.
- ²⁸ U.S. Bureau of Economic Analysis, "Advisory committee on data for evidence building," last modified March 23, 2022, <https://www.bea.gov/evidence>.
- ²⁹ U.S. Congress, House, *Foundations for Evidence-Based Policymaking Act of 2018*.
White House Office of Science and Technology Policy, *Increasing access to the results of federally funded scientific research*, by John Holdren, memorandum for the heads of executive departments and agencies (Washington, DC, February 22, 2013), https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/ostp_public_access_memo_2013.pdf.
Equitable Data Working Group, *A vision for equitable data: Recommendations from the Equitable Data Working Group*, April 22, 2022, <https://www.whitehouse.gov/wp-content/uploads/2022/04/eo13985-vision-for-equitable-data.pdf>.
- ³⁰ See, for example, National Science Foundation, "Dear Colleague letter: Effective practices for data," May 20, 2019, <https://www.nsf.gov/pubs/2019/nsf19069/nsf19069.jsp>.
- ³¹ Nick Hart and Kody Carmody, *Barriers to using government data* (Bipartisan Policy Center, October 2018), <https://bipartisanpolicy.org/download/?file=/wp-content/uploads/2018/10/Barriers-to-Using-Government-Data.pdf>.
Amy O'Hara and Carla Medalia, "Data sharing in the federal statistical system: Impediments and possibilities." *The ANNALS of the American Academy of Political and Social Science* 675, no. 1 (2017): 138–150, <https://journals.sagepub.com/doi/full/10.1177/0002716217740863>.
Andrew Reamer, Julia I. Lane, Ian Foster, and David Ellwood, eds., *Developing the basis for secure and accessible data for high impact program management, policy development, and scholarship*, Sage, 2018, <https://www.aapss.org/volumes/developing-the-basis-for-secure-and-accessible-data-for-high-impact-program-management-policy-development-and-scholarship/>.
- ³² O'Hara and Medalia, "Data sharing in the federal statistical system: Impediments and possibilities," 138–150.
- ³³ U.S. Office of Personnel Management, Chief Human Capital Officers Council, *Position Classification Flysheet for Data Science Series, 1560* (December 2021), <https://chcoc.gov/sites/default/files/Flysheet-Data-Science-Series-Final.pdf>.
- ³⁴ O'Hara and Medalia, "Data sharing in the federal statistical system: Impediments and possibilities." 138–150.
- ³⁵ Michelle M. Mello et al., "Waiting for data: Barriers to executing data use agreements." *Science* 367, no. 6474 (2020): 150–152, <https://www.science.org/doi/full/10.1126/science.aaz7028>.
- ³⁶ "Statistical programs & standards," Executive Office of the President, Office of Management and Budget, accessed April 27, 2022, <https://www.whitehouse.gov/omb/information-regulatory-affairs/statistical-programs-standards/>.

The statistical agencies are the Bureau of Economic Analysis (Department of Commerce, DOC), the Bureau of Justice Statistics (Department of Justice, DOJ), the Bureau of Labor Statistics (Department of Labor), the Bureau of Transportation Statistics (Department of Transportation, DOT), the Census Bureau (DOC), the Economic Research Service (United States Department of Agriculture, USDA), the Energy Information Administration (DOE), the National Agricultural Statistics Service (USDA), the National Center for Education Statistics (Department of Education), the National Center for Health Statistics (United States Department of Health and Human Services, HHS), the National Center for Science and Engineering Statistics (NSF), the Office of Research, Evaluation, and Statistics (Social Security Administration), the Statistics of Income Division (United States of Transportation, DOT), the Microeconomic Surveys Unit (Federal Reserve System), the Center for Behavioral Health Statistics and Quality, the Substance Abuse and Mental Health Services Administration (HHS), and the National Animal Health Monitoring System and Animal and Plant Health Inspection Service (USDA).

Source: Executive Office of the President, Office of Management and Budget, "Statistical Policy Directive No. 1: Fundamental Responsibilities of Federal Statistical Agencies and Recognized Statistical Units," *Federal Register* 79, no. 231 (December 2, 2014): 71610, <https://www.govinfo.gov/content/pkg/FR-2014-12-02/pdf/2014-28326.pdf>.

- ³⁷ Potentially including the Temporary Assistance for Needy Families (TANF) program, the Supplemental Nutrition Assistance Program (SNAP), the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC), and the Centers for Disease Control and Prevention (CDC).
- ³⁸ Baker et al., *Workshop Report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence*.
- ³⁹ "ICESat-2," National Aeronautics and Space Administration, National Snow and Ice Data Center, accessed April 27, 2022, <https://nsidc.org/data/icesat-2>.
- ⁴⁰ "2021 action plan: Agency actions," Executive Office of the President, Office of Management and Budget, Federal Data Strategy, accessed April 27, 2022, <https://strategy.data.gov/2021/action-plan/agency-actions>.
- ⁴¹ "Office of e-government & information technology," Executive Office of the President, Office of Management and Budget, accessed April 27, 2022, <https://www.whitehouse.gov/omb/management/egov/>.
- ⁴² The vision of the Council is to "to improve government mission achievement and increase the benefits to the Nation through improvement in the management, use, protection, dissemination, and generation of data in government decision-making and operations." Source: General Service Administration and Office of Management and Budget, Federal CDO Council, "Home," accessed April 27, 2022, <https://www.cdo.gov/>.
- ⁴³ "PyTorch," PyTorch, accessed April 28, 2022, <https://pytorch.org/>.
- ⁴⁴ "TensorFlow," TensorFlow, accessed April 28, 2022, <https://www.tensorflow.org/>.
- ⁴⁵ "MLCommons," MLCommons, accessed April 28, 2022, <https://mlcommons.org/en/>.
- ⁴⁶ "What is an Edge Database, and why do you need one?" Objectbox, accessed April 27, 2022, <https://objectbox.io/what-is-an-edge-database-and-why-do-you-need-one/>.
- ⁴⁷ "What we do," XSEDE, accessed April 27, 2022, <https://www.xsede.org/about/what-we-do>.
- ⁴⁸ "TensorFlow."
- ⁴⁹ "PyTorch."
- ⁵⁰ "Apache Spark," Apache Spark, Apache Software Foundation, accessed April 28, 2022, <https://spark.apache.org/>.
- ⁵¹ "Basic local alignment search tool," National Institutes of Health, National Library of Medicine, National Center for Biotechnology Information, accessed April 27, 2022, <https://blast.ncbi.nlm.nih.gov/Blast.cgi>.
- ⁵² "SAGE: A software-defined sensor network," Argonne National Laboratory, accessed April 27, 2022, <https://www.anl.gov/mcs/sage-a-softwaredefined-sensor-network>.

-
- ⁵³ Gil and Selman, *A 20-year community roadmap for artificial intelligence research in the U.S.*
Helen Wright, "Robotics Roadmap for US Robotics: From Internet to Robotics, 2020 Edition," CCC Blog, September 9, 2020, <https://cccblog.org/2020/09/09/robotics-roadmap-for-us-robotics-from-internet-to-robotics-2020-edition/>.
- ⁵⁴ See "AI R&D testbed inventory," NITRD, last updated December 17, 2021, <https://www.nitrd.gov/apps/ai-rd-testbed-inventory/>.
- ⁵⁵ "MLCommons."
- ⁵⁶ "The Grand Challenge," DARPA, accessed April 28, 2022, <https://www.darpa.mil/about-us/timeline/-grand-challenge-for-autonomous-vehicles>.
- ⁵⁷ "Robotarium," Robotarium, The Robotarium Project, Georgia Institute of Technology, accessed April 28, 2022, <https://www.robotarium.gatech.edu/>.
- ⁵⁸ "Robotics test facility," NIST, accessed April 28, 2022, <https://www.nist.gov/laboratories/tools-instruments/robotics-test-facility>.
- ⁵⁹ "POWDER-RENEW," Platforms for Advanced Wireless Research, accessed April 28, 2022, <https://advancedwireless.org/salt-lake-city/learn-more-powder-renew/>.
- ⁶⁰ "The Truman platform," Cornell Social Media Lab, Cornell University, accessed April 28, 2022, <https://socialmedialab.cornell.edu/the-truman-platform/>.
- ⁶¹ "Face recognition vendor test (FRVT)," NIST, accessed April 28, 2022, <https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt>.
- ⁶² "COSMOS," COSMOS Testbed Main Site, accessed April 28, 2022, <https://www.cosmos-lab.org/>.
- ⁶³ "Energy Smart Data Center," Pacific Northwest National Laboratory, last updated September 2014, <https://www.pnnl.gov/computing/resources/esdc/>.
- ⁶⁴ "CyberNET testbed," Pacific Northwest National Laboratory, accessed April 28, 2022, <https://www.pnnl.gov/projects/cybernet-testbed>.
- ⁶⁵ "What is Dioptra?" Dioptra 0.0.0 documentation, NIST, accessed April 28, 2022, <https://pages.nist.gov/dioptra/>.
- ⁶⁶ "National Transportation Research Center," Oak Ridge National Laboratory, accessed April 28, 2022, <https://www.ornl.gov/facility/ntrc>.
- ⁶⁷ "About AERPAAW," AERPAAW, accessed April 28, 2022, <https://aerpaw.org/>.
- ⁶⁸ Junyang Lin et al., "M6: A Chinese multimodal pretrainer," *arXiv preprint arXiv:2103.00823* (2021), <https://doi.org/10.48550/arXiv.2103.00823>.
- ⁶⁹ "Kaggle," Kaggle Inc., accessed April 28, 2022, <https://www.kaggle.com/>.
- ⁷⁰ "The Grand Challenge," DARPA.
- ⁷¹ "Robocup," accessed May 23, 2022, <https://www.robocup.org/>.
- ⁷² "Overview," Center for Digital Agriculture, University of Illinois, accessed April 28, 2022, <https://digitalag.illinois.edu/autonomous-farm/overview/>.
- ⁷³ Austin Harris, Jose Stovall, and Mina Sartipi, "MLK Smart Corridor: An urban testbed for smart city applications," *2019 IEEE International Conference on Big Data (Big Data)*, 3506–11, 2019, <https://doi.org/10.1109/BigData47090.2019.9006382>.
- ⁷⁴ "Digital city testbed center," Portland State University, accessed April 28, 2022, <https://www.pdx.edu/digital-city/>.
- ⁷⁵ "Smart community testbed," Virginia Innovation Partnership Corporation, accessed April 27, 2022, <https://www.virginiaipc.org/smart-community-testbed>.
- ⁷⁶ "Robotarium," Georgia Institute of Technology.
- ⁷⁷ "OSG," The OSG Consortium, accessed April 28, 2022, <https://opensciencegrid.org/>.
- ⁷⁸ "Cloudbank," Cloudbank, accessed April 28, 2022, <https://www.cloudbank.org>.

-
- ⁷⁹ "The Big Data Innovation Hubs," The Big Data Innovation Hubs, accessed April 28, 2022, <https://bigdatahubs.org/>.
- ⁸⁰ "OSG."
- ⁸¹ "Training overview," XSEDE User Portal, accessed April 27, 2022, <https://portal.xsede.org/training/overview>.
- ⁸² "Technical resources," Cloudbank, accessed April 27, 2022, <https://cloudbank-project.github.io/cb-resources/>.
- ⁸³ "Berkeley Data Stack," accessed April 27, 2022, <https://data.berkeley.edu/academics/campus-resources/berkeley-data-stack>.
- ⁸⁴ Usable cybersecurity involves the consideration of human factors in the context of cybersecurity (for example, how people make security decisions, their perceptions of security, how they interact with security mechanisms, and the associated challenges of those interactions) with a goal of designing cybersecurity technologies, processes, and policies that are both usable and result in improved security outcomes.
- ⁸⁵ "About," Trusted CI, accessed April 28, 2022, <https://www.trustedci.org/framework>.
- ⁸⁶ Laura Taylor, "FedRAMP: History and future direction," *IEEE Cloud Computing* 1, 3 (2014): 10-14; see also <https://fedramp.gov>.
- ⁸⁷ OMB Memorandum M-22-09, January 2022, <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.
Executive Order 14028, May 2021, <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>.
- ⁸⁸ "What is the Five Safes framework?" UK Data Service, accessed April 28, 2022, <https://ukdataservice.ac.uk/help/secure-lab/what-is-the-five-safes-framework/>.
- ⁸⁹ "Five Safes framework," Coleridge Initiative, accessed April 28, 2022, <https://coleridgeinitiative.org/adrf/five-safes/>.
"Five Safes: Designing data access for research," University of the West of England Bristol, Economics Working Paper Series, 1601, 2016, https://www.researchgate.net/publication/292975549_Five_Safes_designing_data_access_for_research.
- ⁹⁰ Presentations to the TF; J. Lane, V. Stodden, S. Bender, & H. Nissenbaum, (Eds.). (2014). *Privacy, big data, and the public good: Frameworks for engagement*. Cambridge University Press.
- ⁹¹ Jeffrey Dastin, "Amazon scraps secret AI recruiting tool that showed bias against women," Reuters, October 10, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.
- ⁹² Steve Lohr, "Economists pin more blame on tech for rising inequality," *New York Times*, updated January 20, 2022, <https://www.nytimes.com/2022/01/11/technology/income-inequality-technology.html>.
- ⁹³ Sammy Westfall, "U.N. official calls for moratorium on artificial intelligence tools that breach human rights," *Washington Post*, September 15, 2021, <https://www.washingtonpost.com/world/2021/09/15/un-ai-moratorium/>.
- ⁹⁴ "Public Input on a National AI research resource implementation plan," National Artificial Intelligence Initiative, accessed April 28, 2022, <https://www.ai.gov/nairrtf/86-fr-39081-responses/>.
- ⁹⁵ "Responsible computing research: Ethics and governance of computing research and its applications," National Academies, National Academy of Sciences, accessed April 28, 2022, <https://www.nationalacademies.org/our-work/responsible-computing-research-ethics-and-governance-of-computing-research-and-its-applications>.
- ⁹⁶ "Federal plain language guidelines," Plainlanguage.gov, Plain Language Action and Information Network, accessed April 28, 2022, <https://www.plainlanguage.gov/guidelines/>.
- ⁹⁷ Dr. Parashar was named TF co-chair in October 2021.

