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

Response to the NAIRR Request for Information

Submitting Organizations:

AI Institutes established in 2020:
- [NSF AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography](#)
- [NSF AI Institute for Foundations of Machine Learning](#)
- [NSF AI Institute for Student-AI Teaming](#)
- [NSF AI Institute for Molecular Discovery, Synthetic Strategy, and Manufacturing](#)
- [NSF AI Institute for Artificial Intelligence and Fundamental Interactions](#)
- [USDA-NIFA AI Institute for Next Generation Food Systems](#)
- [USDA-NIFA AI Institute for Future Agricultural Resilience, Management, and Sustainability](#)

AI Institutes established in 2021:
- [NSF AI Institute for Collaborative Assistance and Responsive Interaction for Networked Groups](#)
- [NSF AI Institute for Advances in Optimization](#)
- [NSF AI Institute for Learning-Enabled Optimization at Scale](#)
- [NSF AI Institute for Intelligent Cyberinfrastructure with Computational Learning in the Environment](#)
- [NSF AI Institute for Future Edge Networks and Distributed Intelligence](#)
- [NSF AI Institute for Edge Computing Leveraging Next-generation Networks](#)
- [NSF AI Institute for Dynamic Systems](#)
- [NSF AI Institute for Engaged Learning](#)
- NSF AI Institute for Adult Learning and Online Education
- [The USDA-NIFA Institute for Agricultural AI for Transforming Workforce and Decision Support](#)
- [The USDA-NIFA AI Institute for Resilient Agriculture](#)

Background: The National AI Institutes Program established **18 AI institutes in 2020 and 2021, with a total investment of $360M over 5 years**. Each institute covers both foundational and applied (or "use-inspired") AI research, and they collectively address the use of AI for a wide range of societal benefits. These institutes collectively support **over 500 researchers in over 40 states**, in a broad range of areas of core AI research, including foundational machine learning, scientific machine learning, autonomous systems, computer vision, natural language processing, multimodal interaction, edge intelligence, optimization, efficient AI, trustworthy AI, interpretable AI, fair/unbiased AI, ethics in AI, and AI education, as well as numerous broad application domains, such as agriculture, computer system design, education and adult learning, engineered systems, home healthcare, molecular manufacturing, nutrition, physics, and weather and oceanography. Over the next five years and probably much longer, the AI-focused research, education, and outreach needs of these 18 institutes will represent a very large and highly impactful user community for the NAIRR. The responses below represent the combined needs of these highly multidisciplinary AI Institutes.

Most importantly, the intensive planning and discussions that went into creating these 18 institutes have uncovered numerous important requirements for multidisciplinary AI research, some (but not all) of which are summarized briefly in the responses below. Our overarching process recommendation is to consult a representative set of individuals from these 18 institutes in the planning and creation of the NAIRR.

Q1. What options should the Task Force consider for any of roadmap elements A through I, and why?

- **(Item A) Goals:** We believe that the NAIRR should consider the needs of both foundational and multidisciplinary use-inspired AI research for a wide range of target domains. It should ensure that competitive processes for access to these resources are folded into competitive processes for research funding so that funded research teams are guaranteed access to the computational resources they need at the same time they are awarded other resources. It should consider all tiers of resources, spanning from high-end training to low-end edge devices and everything in between.

  (Item A) Metrics: Metrics of success should span from the foundational to the applied. Some important metrics include (a) measures of cost, e.g., in comparison with publicly accessible commercial alternatives; (b) measures of access, spanning race/ethnicity/gender, career stages, university type; (c) measures of data set availability and usefulness, e.g., through scientifically valid surveys of various potential user communities; (d) measures of research successes, such as important research breakthroughs, standards adopted, new collaborations, etc.; (e) measures of performance benefits, e.g., on standard AI benchmarks or in quantifiable use-inspired applications; (f) measures of commercial impact, e.g., technology transfer, licensing, or startups that are spun off; and (g) qualitative results of an annual review process, conducted by a respected team of independent external evaluators.

- **(Item B) Responsible agency or organization:** We would recommend NSF as the main agency to create and govern this resource. Their mission addresses a very broad range of foundational and applied scientific research, they already fund large national computational resources, and they have previously coordinated similarly broad programs.

- **(Item C) Governance, Oversight, Resource Allocation:**
  - The governing body should include a three-way public-private partnership of academic leaders, representatives from federal funding agencies, and industry. A non-profit entity should serve as the coordinating entity, with oversight by the combined set of above stakeholders.
  - The NAIRR governing body should include subgroups with expertise in important subareas of AI as well as key application domains. The NAIRR should leverage the entire research and development community, by enabling community contributions of software

tools, curated data sets, tutorials and usage manuals, discussion forums, evaluation studies, etc., all supervised by appropriate NAIRR experts.
  - Resource allocation processes should not separate research funding (e.g. NSF grants to support PIs and other personnel in the conduct of research projects) from the computational resources needed to undertake that research. Access to the NAIRR should be allocated as part of the ordinary funding process for research programs and without requiring a separate proposal.
  - For those without existing research funding or proposals in review, an open process similar to XSEDE proposals could be used, but with much more flexibility, e.g., without an expectation of predictable milestones or a timeline of outcomes, and also without requiring significant existing HPC experience, which blocks new research users.
- **(Item D) Capabilities:** We recommend a number of specific capabilities for the NAIRR:
  - **High scalability** in both compute capacity and memory size, supporting both traditional parallelism e.g., using MPI/OpenMP, and heterogeneous parallelism for machine learning, e.g., GPUs and deep learning (DL) accelerators.
  - **Configurable edge computing** ecosystems that enable experimentation with a wide range of edge AI applications, spanning very low-end systems (e.g., in IoT devices) to intermediate scale (e.g., mobile computing or AR/VR/XR) to higher-end systems (e.g., autonomous vehicles or smart cities or smart manufacturing).
  - **Data sets and data management workflows:** Curated high-quality labeled and unlabeled data sets spanning as many application domains as possible, as well as an easy-to-use pipeline for data cleaning, curation, imputation, and fusion. Outreach and awareness (along with sustained support) to identify, disseminate and grow these data sets.
  - **Benchmarking:** Exclusive and/or isolated access to systems or subsystems for benchmarking.
  - **Future-proofing:** Increasing scalability requirements over time as community needs and use cases expand.
  - **Training materials:** Extensive training materials for novices and experts alike, including training workshops, challenge competitions, and curriculum development.
  - **User support and ease of use:** Compute resources that are easy to access, understand, configure and use. Technical support for users who need assistance with getting started using the NAIRR, or with using it more efficiently and effectively.
- **(Item E) Barriers to dissemination and use:** *No response provided.*
- **(Item F) Security requirements:** CUI (Controlled Unclassified Information) data access policies, e.g., for energy-domain data, HIPAA and FERPA protected data sets, proprietary commercial and end-user data sets, secure storage and processing of proprietary or sensitive data.

- **(Item G) Anonymity, privacy, and accountability:** End-users should enjoy strong privacy and intellectual property (IP) protection guarantees for their identity, activities, results, and outcomes. Team members *and* core community contributors (selected by a well-specified public process) should be publicly identified and accountable for their contributions. NAIRR representatives should be available to answer questions about policies and practices.
- (Item H) As emphasized in the response to Q5, below, the private sector should be encouraged and even expected to play a major role in funding the creation, maintenance, and long-term sustainability of the NAIRR, and in tracking and acknowledging the commercial benefits that accrue from it.

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

The most important capabilities and services the NAIRR should prioritize are:

- **Personnel funding:** AI infrastructure and data sets are not just about the "big iron": there is a large "people expense" and the NAIRR should be funded to support this component of the expense.
- **Datasets:** Dedicated funding for creating/curating/annotating/maintaining datasets is critical.
- **Training and user support**: Training -- including online, remote tutorials -- for use of the computational and data resources is essential. So is a technical support team dedicated to assisting both novice and advanced users to make best use of the resources.
- **System heterogeneity:** The computational infrastructure should include a diverse range of hardware, software, and networking, e.g., CPUs, GPUs, FPGAs, and accelerators, spanning high-end to low-end systems.
- **Scholarships:** Student funding for diversifying the user base of the NAIRR will be important.

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

- **Ensure a diverse governing body:** Ensure that the group and agency governing this resource is diverse in all respects (not just race/ethnicity/gender but across university type, career stages, etc). This will bring a wealth of ideas to ensure a level playing field for all in terms of access.
- **Develop and publish an Ethics Policy:** The starting point is establishing a culture of ethics. This can be expressed in a clear policy statement that defines expectations for the ethical and responsible use of AI, enforceable for projects that make use of the NAIRR. Include requirements about the contents of data sets, with a particularly high bar for data that could become part of a formal or de facto standard or otherwise have potential to significantly impact equity and justice.

- **Use equitable and transparent allocation policies:** A traditional, peer-reviewed, proposal-driven system will tend to favor established researchers. If there is a proposal-driven process, it should include mechanisms (e.g., a "junior researcher" track or a track for non-R1 institutions) that allocate a defined fraction of the resources with substantial average allocations.
- **Provide incentives:** Explicitly reserve a fraction of the resources for research programs that focus on enhancing ethical and responsible use of AI. Other funding programs that use the NAIRR should be required to encourage these principles, and provide supplementary support for projects that successfully demonstrate that they have adhered to these goals.
- **Insulate researchers from pressure:** Excessive reliance on *direct* industry funding can undermine important ethics-related research and education in AI. Impartial, federal funding sources such as NSF are crucial. Industry contributions should be channeled through federal agencies with carefully constrained industry participation.
- **Require adherence:** The NAIRR should also require that (many or most) projects using the NAIRR include an explicit component to measure and/or train people in ethical and responsible use of AI.
- **Include a strong ombudsman:** Fund and support an independent ombudsman with authority to investigate and penalize projects that violate the Ethics Policy.
- **Be accountable:** Measure and publicize key metrics regarding the NAIRR's user community, including equity, bias, diversity, and other important goals.

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

- **Existing large-scale systems:** Two obvious examples are the high-end computing resources of the NSF and DOE-funded supercomputer centers (e.g., Delta, Cheyenne, Stampede2, Frontera) and private cloud services (Google Cloud, AWS, Microsoft Cloud). Learning from all of this is highly recommended but you should also see what has worked and not worked well, through surveys of the users. Some of these approaches do not scale well and do not provide sufficient resources to actually achieve success while others do.
- **Large testbeds for evaluation:** Large existing research testbeds can provide a strong basis for experimental evaluation of AI-driven modeling and prediction techniques in diverse application domains. Some examples include the four city-scale research testbeds for advanced wireless research funded by the $100M PAWR project; NIST laboratory testbeds ranging from manufacturing robotics to software tools to fire protection and others; the SoyFACE facility for climate resilience in agricultural crops; the ESnet 100G SDN testbed for evaluating advanced networking research; NSSL's Hazardous Weather Testbed and National Weather Radar Testbed; and numerous others.

- **Data sets from large public testbeds:** In addition to experimental evaluation, the above testbeds may also be valuable sources of labeled and unlabeled data sets for machine learning research.
- **Commercial data sets:** Private sector companies have vast proprietary datasets and are usually unwilling to share them, but may be incentivized to share representative samples, e.g., with some NAIRR funding allocated for this purpose.
- **Centralizing coordination of existing data sets:** Many efforts, e.g., federally funded projects, generate valuable data sets and have a requirement (or a goal) to make these public, but there are few central directories or organizational structures for coordinating and supporting access to these resources (the National Data Service may be one example). One component of NAIRR funding could be allocated to publicizing the need for such data sets, and to identifying, curating, coordinating, and supporting access to related datasets from disparate sources.

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

US and international companies have benefited tremendously from international research efforts in artificial intelligence, and they should be strongly encouraged to play a major role in funding the NAIRR. Many of the most promising applications of AI are in large, highly profitable, and fast-growing industry sectors that are likely to see major gains in new technological capabilities, new product categories, and large new revenue streams. *The NAIRR will play a critical role in enabling research into such applications over the next decade or more. Large, highly profitable companies today pay virtually nothing for these benefits, and that situation must change*. Moreover, part of the reason for this situation is the lack of careful accounting of the financial impact of past research, and that should be rectified as well. **Some ways in which companies could, and *should*, contribute to the NAIRR include:**

- **Pay for it:** Substantial industry contributions to joint federal+private funding programs to support the NAIRR. For example, four companies contributed $5M each to five national AI institutes in the 2021 round, representing 20% of the total funding for these five institutes. Other NSF-funded programs, such as PAWR, RINGS, MLWins etc., have also involved public-private partnerships.
- **Donate to it:** No-cost donation of storage and/or compute time on large cloud-based services and edge-compute hardware and networking for urban and rural applications, including access to both general-purpose and specialized hardware. Private entities could also contribute in-kind resources, or provide access to production systems or valuable datasets, through "grand challenge" competitions.
- **Create data sets for it:** Creation and publication of non-proprietary data sets that are representative of commercial workloads and assistance to researchers in accessing and using these data sets.

- **Support and diversify the students using it:** Scholarship and internship programs that expand access and skills among underserved groups.
- **Encourage entrepreneurial outcomes from it:** Venture capital to incentivize spinoff companies that emerge from NAIRR-enabled research.
- **Track and publicly acknowledge new revenue streams that accrue from it:** Companies should explicitly account for and publicly acknowledge new technologies, new product categories, new services, and new workforce training enabled by NAIRR, including a public summary of the revenues and profits that accrue from these new capabilities.

Q6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

There are roughly 3,900 accredited colleges and universities in the U.S. Democratized access requires an open infrastructure, along with federation of existing agencies' assets with this infrastructure. Some suggestions for democratizing access to the NAIRR are:

- The NAIRR should design its funding structure and competitive grant programs to spread the resources across a wide range of educational institutions, along several dimensions: research capacity (e.g., R1 research universities, 4-year colleges, and community colleges), demographics, geographic location, etc.
- One key obstacle for smaller institutions is lack of research funding. Access to NAIRR should *not* require existing research funding. On the contrary, it should provide explicit access to individuals without such funding who can use the system to bootstrap their research programs.
- Research proposals for federal funding programs that make use of the NAIRR should be required to include a strong BPC component focused on enabling training and access to the NAIRR.
- The NAIRR can provide automatic access to peer-reviewed seedling programs, like the NSF CISE Research Initiation awards, that are limited to junior researchers who have not yet secured independent federal funding.
- All publications describing research that uses the NAIRR should be made freely available to the community, without barriers of paid subscriptions.

NAIRR and RFI Links

- **NAIRRTF: THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH RESOURCE TASK FORCE**
- **Request for Information (RFI) on NAIRR Implementation Plan**

Questions for Response

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

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

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

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

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

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