

MEETING SUMMARY

National Artificial Intelligence Research Resource Task Force

Meeting #4

December 13, 2021

Meeting Summary

The fourth meeting of the National Artificial Intelligence Research Resource Task Force (NAIRR TF) was held online via Zoom on December 13, 2021, 11:00 AM–6:00 PM EST.

Welcome and Administrative Remarks

The meeting started at 11:09 AM EST.

Dr. Lynne Parker (White House Office of Science and Technology Policy), NAIRR TF Co-Chair, opened the meeting. Dr. Parker introduced Dr. Manish Parashar (Office Director, Office of Advanced Cyberinfrastructure, National Science Foundation), who had transitioned into the role of Co-Chair previously held by Dr. Erwin Gianchandani. (Dr. Gianchandani, previously the Deputy Assistant Director for Computer and Information Science and Engineering at the National Science Foundation, transitioned to a new role as Senior Advisor to the Director for Translation, Innovation, and Partnerships, which prompted this transition.)

Dr. Parashar entertained a motion to approve the minutes from the prior NAIRR TF meeting; the motion passed.

Dr. Parashar then introduced the agenda, detailing the primary topics that would be taken up during the meeting.

The session ended at 11:19 AM EST.

Overarching NAIRR Vision

The session started at 11:19 AM EST.

Dr. Parker presented an overarching vision for the NAIRR, which was developed by the Co-Chairs in collaboration with TF members, who provided input during individual meetings. Dr. Parker laid out areas where the TF members appear to have reached consensus, including:

- The NAIRR's objectives are to strengthen the Nation's AI innovation ecosystem and, by extension, economic growth and competitiveness, by lowering barriers to entry in the field, promoting AI skills and knowledge for developing the AI workforce, broadening participation to include all segments of the Nation, and supporting innovative AI-relevant research;
- The NAIRR should be a federated ecosystem of existing and new resources;
- The resource should strive to achieve the attributes of transparency, trust, security, robustness, accessibility, independence, scalable functionality, sustainability, oversight, and accountability;
- The baseline user profiles include AI researchers, researchers conducting use-inspired AI research and using AI to advance other fields, and students;

- The NAIRR should be open to U.S.-based AI researchers and students at academic institutions, non-profit research organizations, national laboratories, federally-funded research and development centers, and startup companies or research organizations that have been awarded federal Small Business Innovation Research (SBIR) or Small Business Technology Transfer (STTR) grants;
- The openness of research conducted using the NAIRR should be consistent with the policies of the sponsoring federal agencies; and
- Access to all of the NAIRR resources should be subject to clear use policies and user agreements.

Dr. Parker then walked the TF through additional considerations requiring further discussion, notably the user base and allocation of resources, including whether and under what circumstances private-sector researchers might access the NAIRR. TF members discussed these additional considerations, noting that it is important to encourage interdisciplinary AI work, such as in biology, manufacturing, and agriculture, and that some big companies that are not digitally native may have a need for AI resources. TF members discussed how it will be necessary to consider review processes to determine public benefit of private-sector usage of the resource; how private-sector users might pay fees if they want to keep the work proprietary; and how to partner with diverse federal agencies to support an interdisciplinary user base. TF members also identified the need to decrease review process times, potentially by using a tiered review approach, as well as the challenges inherent in reviewing for a multidisciplinary portfolio involving multiple agencies.

Dr. Parker concluded the discussion and noted that additional topics on the NAIRR vision should be raised to the Co-Chairs by January 12.

The session ended at 12:00 PM EST.

Readout and Discussion of Draft Recommendations: User Resources Working Group

The session started at 12:00 PM EST.

Dr. Parker discussed how TF input on Working Group (WG) recommendations would be incorporated. She described how, since the last meeting, three WGs had met to develop proposed conclusions and recommendations for consideration by the TF. Additionally, she stated that the recommendations would lay the foundation for the interim report to be released in the spring. If no consensus is reached, the concerned and interested parties would be convened to work out a solution, which would be presented at the next TF meeting in February.

Dr. Fei-Fei Li presented the findings and recommendations of the User Resources WG. Topics included NAIRR system design criteria, including integration of services, resources, data, and training material; user-centered design; scalability and extensibility of data and compute; speed; user training and education; and tiered support. The WG recommended that the NAIRR leverage user portal concepts from existing state-of-the-art approaches, explore the possibility of outsourcing portal development, and identify and curate appropriate education and training materials for different skill levels, to include those available from NAIRR resource providers.

Discussion also touched upon how to support users with different experience levels, from novice to advanced, and across different research domains. The NAIRR could act as a liaison to support community building among users, for example, through forums, and staff could support engagement

with providers and Federal experts. Relationships with existing Centers of Excellence and different existing technical solutions to support collaborative data science could also be leveraged.

The session ended at 12:49 PM EST.

Readout and Discussion of Draft Recommendations: Testbed/Testing Resources Working Group

The session started at 12:49 PM EST.

Dr. Andrew Moore presented the recommendations of the Testbed/Testing Resources WG. The WG had deliberated on whether and how the NAIRR should spend resources on testbeds. The WG recommended that the NAIRR should spend between 5% to at most 10% of its resources on testbeds. A NAIRR Testbed Office could maintain a world-class catalogue of testbeds with pointers to those resources and develop some testbeds that would not otherwise be available. The WG recommended that the NAIRR provide infrastructure to support Open Book Modeling, Closed Book Modeling, and Simulated Perceive-Decide-Act competitions, but de-emphasize Real-World Perceive-Decide-Act competitions.

TF members discussed the definitions of benchmarking and testbeds, including whether the NAIRR would be a place for hardware comparisons. Members also raised the question of the role of living laboratories, agreeing that creation of such resources would likely be outside the NAIRR's scope but acknowledging the NAIRR could connect to existing living labs and make their datasets available. The TF also discussed whether testbeds would be bringing their own computational and data resources, how the NAIRR could support the creation of a knowledge repository or play a role in supporting researchers who want to create testbeds, and the importance of ensuring that the provisioning of urgently-needed computational resources remains a NAIRR priority.

The session ended at 1:38 PM EST.

Break: 1:38-2:20PM EST

Readout and Discussion of Draft Recommendations: Data Working Group

The session started at 2:20 PM EST.

Drs. Daniela Braga and Julia Lane presented the recommendations of the Data WG. The WG concluded that the full value of AI is often not realized without high-quality, trusted, dense, and transparent data. The WG suggested that, in addition to open data, statistical and administrative government data and data generated by federally-funded research should also be made available through the NAIRR. The WG recommended that the NAIRR coordinate a network of trusted data/computational providers and hosts for a transparent and responsible data marketplace that incentivizes the provisioning of high-quality data, accessible via a search-and-discovery platform. The WG further recommended a tiered-access approach through the NAIRR portal, leveraging the "Five Safes" framework for decision making about data access and use. Data governance policies and practices should be established and periodically updated, and the NAIRR should provide funding to independent oversight entities and mechanisms for public engagement to support transparency and reduce the potential for harms. Substantial resources would also be needed for infrastructure, technical support staff who can help with data curation, and training programs on NAIRR data policies and acceptable practices. TF members then discussed the challenges of creating an

incentive structure for data holders and a need for repositories of natural-language, human-speech, video, and high-quality, tagged data. The challenge of providing export control for datasets that should not be downloaded to a separate machine was also discussed as was the fact that even if sensitive data constitute only a small percentage of the data made available through NAIRR, such features will take significant time and effort to implement and monitor.

The session ended at 3:21 PM EST.

Panel: Privacy, Civil rights, and Civil Liberties Requirements

The session started at 3:21 PM EST.

The panel comprised the following individuals:

- 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; Founder/President, Data & Society;
- 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; and
- Deborah Raji, Fellow, Mozilla Foundation;

Dr. Parashar introduced the panel, and each panelist spoke for approximately five minutes, discussing concerns associated with AI and machine learning, including:

- Issues might manifest both in the conduct of research and in the downstream effect of research findings;
- Consent given to participate in research by a minority of a population could impact the broader population;
- AI can perpetuate an inequitable status quo;
- Biases can arise as a result of following the suggestions of algorithms;
- New uses of machine learning and AI are likely to include unforeseen issues related to trustworthiness and fairness; and
- Machine-learning models may show bias when they are attacked.

Panelists proposed suggestions for NAIRR governance to address these challenges, including:

- Bring civil society stakeholders in early and embed them throughout the processes of the NAIRR;
- Design the NAIRR to prevent abuse and manipulation;
- Develop protocols, communication infrastructure, and an engagement plan for use in the event research that used NAIRR resources causes harm;
- Develop strategies to protect research subjects and populations they represent, which could be done by creating a NAIRR Audit and Oversight Board with diverse representation and representation from civil society; and
- Design review and feedback mechanisms into the NAIRR governance to ensure the NAIRR is supporting more equitable access and use.

Panelists also discussed the need not only for diversity in disciplines (e.g., lawyer, sociologist, philosopher, privacy expert) but also in demographics of those developing, designing, auditing, and

executing AI. Panelists suggested providing training and resources intended specifically for community colleges and high-school students, and including researchers from Historically Black Colleges and Universities, Hispanic-Serving Institutions, other Minority Serving Institutions, and civil society originations.

Panelists suggested that the NAIRR concentrate on making it as easy as possible for researchers who focus on fairness to examine and build upon the results of researchers who focus on specific applications of machine learning. Additionally, the NAIRR could enable building tools and models that support third-party auditing and oversight, along with appropriate auditing and accountability through carefully-designed Terms of Service.

Following remarks from each panelist, Dr. Parashar moderated a discussion with TF members. TF members contemplated mechanisms to deliberate and decide on the approval of individual research projects on the NAIRR, a requirement that a body involved in this be representative of the community that could be affected by the research results, and ways for the NAIRR to address the concerns of civil society. Panelists suggested establishing guidelines for minimum standards of model performance or due diligence; creating a framework for assessing ethical AI; and considering the harmful experiences shared by representatives of communities when thinking through potential algorithmic harms.

The session ended at 4:32 PM EST.

Break: 4:32-4:40 PM EST

Briefing: Security and Access Control Considerations for the NAIRR

The session started at 4:40 PM EST.

Dr. Emily Grumbling and Morgan Livingston (Science and Technology Policy Institute) provided an overview of key issues related to NAIRR security requirements. They noted that the NAIRR will be subject to standard cybersecurity threats as well as AI/machine learning-specific threats. The NAIRR will require security policies, staff, and governance entities as well as implementation of user access controls. A least-privilege, tiered access approach with either role-based or attribute-based access controls would help to minimize the risk of misuse, leak, compromise, corruption, or theft of system resources. While privacy- and security-preserving machine-learning methods are unlikely to be universal solutions for protecting the confidentiality of training and testing data, the NAIRR system could help to enable research to advance these and other methods.

The speakers also described how a federated NAIRR model would both increase complexity and provide opportunities to distribute security responsibilities across the partners, noting the TF may wish to identify clear security roles, responsibilities, and requirements for NAIRR partners. They suggested it may be helpful to initiate the NAIRR first as a small pilot or series of pilots to test out the security policies, architectures, and access controls, in order to troubleshoot and harden the system before it goes fully live. The TF can leverage established approaches for access control that align with NAIRR resource components and goals.

The session ended at 5:00 PM EST.

Discussion: Usable Security and User Access Controls

The session started at 5:00 PM EST.

Ms. Elham Tabassi moderated a discussion among TF members on the NAIRR's approach to user access controls, unique security requirements, and balancing usability and security. TF members discussed how the NAIRR could build on cybersecurity best practices. The NAIRR should be the gold standard: data have to be clean, vetted, and accurate. TF members discussed the challenges of a federated resource: as complexity grows, threats grow, and security is more difficult. Usable security will involve making sure access controls are designed such that usability is high and complexity is low. TF members also discussed questions of the scope of the NAIRR and its role in supporting research on privacy and security.

The session ended at 5:23 PM EST.

Briefing: Technical Integration

The session started at 5:24 PM EST.

Dr. Michael Norman presented on the technical integration of resources into a NAIRR cyberinfrastructure, including how the integration of AI/machine learning data repositories and edge computing resources represent a new element to the computational infrastructure ecosystem. The NAIRR might incorporate existing resources, such as cloud, campus, high-performance computing clusters, and government data repositories, as well as new NAIRR resources such as machine-learning data depots, NAIRR prototype systems and testbeds, and edge resources. Dr. Norman additionally presented goals for and examples of technical integration of computational, data, edge, allocations and usage reporting, along with training resources.

TF members discussed the lack of standards for the edge ecosystem and raised the question of whether NAIRR should play a role in establishing standards. Both integrating edge resources and data pose challenges. Members also commented on the need to provide users access to usage data on NAIRR datasets.

The session ended at 5:53 PM EST.

Questions from Public and Meeting Close

The session started at 5:53 PM EST.

Dr. Parker addressed questions submitted by attendees via the webinar Q&A portal. All questions had been answered in written format, except one that was answered live. Questions addressed included how the NAIRR might interact with existing federal enterprise architectures, support technologies such as cloud templates, and support existing businesses certified by Federal agencies.

Dr. Parashar concluded the meeting, thanking members of the TF and the public. He noted that meeting summaries, slide presentations, and details about upcoming meetings can be found at <https://www.ai.gov/nairrtf/>.

Dr. Parashar reminded everyone that the next meeting is scheduled for February 16, 2022, from 11:00AM-5:00PM EST. Details are posted to the Federal Register.

The meeting adjourned at 5:57 PM EST.

Appendix I: Attendance for NAIRR TF Meeting #4

TF Members Present:

Manish Parashar, National Science Foundation (Co-Chair)

Lynne Parker, White House Office of Science and Technology Policy (Co-Chair)

Daniela Braga, DefinedCrowd

Mark Dean, retired (formerly IBM and University of Tennessee, Knoxville)

Oren Etzioni, Allen Institute for AI

Julia Lane, New York University

Fei-Fei Li, Stanford University

Andrew Moore, Google

Michael Norman, University of California, San Diego

Dan Stanzione, University of Texas, Austin

Frederick Streit, Department of Energy

Elham Tabassi, National Institute of Standards and Technology

TF Members Absent:

None