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

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

Via electronic submission at http://www.regulations.gov

Subject: Google Comment Regarding the Office of Science and Technology Policy and National Science Foundation NAIRR Task Force Request for Information

Reference: 86 FR 39081, Document Number 2021-15660

October 1, 2021

Google welcomes the opportunity to provide comments in response to the **"Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource (NAIRR)"** issued jointly by the National Science Foundation and the Office of Science and Technology Policy.

Artificial intelligence (AI) will continue to have a significant, positive impact on society—for example through enabling improvements in healthcare, education, business, and government, and helping to address key challenges like climate change. The implementation and uptake of AI also raises questions about governance, privacy, security, access, economic opportunity, and dual use. We strongly support establishing a NAIRR, and share the government's goal to make AI access more equitable through this resource. The NAIRR is a great opportunity to support increased access for a diverse range of researchers to critical AI and cloud resources such as storage, compute, databases, networking, data analytics, AI services and collaboration tools.

In developing a NAIRR implementation plan, we encourage the Task Force to focus on leveraging the existing and unique capabilities of US academic institutions, government agencies, and industry to further enhance US competitiveness.

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

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

Suggested goal: Provide efficient and expanded access to resources (including storage, compute, databases, networking, data analytics, AI services and collaboration tools) for US AI researchers and practitioners, from all sectors

(including academia; federal, state, and local governments; and private sector), and facilitate cross-sectoral and interdisciplinary projects. Researcher access should be prioritized based on an analysis of the public benefit of the proposed work and cost to provide such benefits. Special consideration should be made for how government and academic researchers may leverage the resource. Metrics for this goal:

- For funded proposals, time between a researcher's request for resources and access to resources
- # of overall NAIRR users, and # of active users
- Demographics of NAIRR users, including race/gender, users from specific sectors, regions, and populations, including government researchers, academics (with disaggregated data on academics from minority serving and emerging research institutions), and startups. We especially encourage use of the NAIRR to facilitate US government AI research, including interagency collaboration on AI and research at national labs.
- # of publications resulting from NAIRR use
- Project impact, measured through some combination of: # and size of follow on projects/funding, qualitative impact analyses, citations, impact factor of journal articles, awards
- # of open-source projects / contributions related to NAIRR use, and use of these contributions by the broader community
- # of interdisciplinary (i.e. led by chemists, biologists, humanities scholars) and cross-sectoral projects and publications, and publication impact
- Applications resulting from NAIRR use, with qualitative evaluation of their impact
- # of unique public/private data sets made available and used by researchers, and quality of de-identification of data in these data sets
- # of machine learning models/architectures made available and used by researchers
- Overall carbon impact and renewable energy savings, as well as carbon impact per project (goal should be carbon neutral to start, and move toward 100% renewable energy use for all compute, with yearly targets for renewable use)

Suggested goal: Workforce development -- train more AI researchers and practitioners, specifically in use of cloud resources. One of the main barriers to successful use of the cloud - as identified by federal programs such as CloudBank and STRIDES -- is that researchers lack training on how to access the cloud, how credits work, and how to appropriately budget for its use. Such training should be built into NAIRR implementation. In addition to training and support, NAIRR administrators should establish a formal process/platform to collect feedback from participating researchers and share with private sector training organizers, where relevant. Metrics for this goal:

- # and features of researchers and research administrators trained in specific aspects of cloud use, and correlation between this training and further use of the NAIRR (i.e. metrics described under the first goal above)
    - Features can include demographics, regional representation, non-AI specialist researchers,
- Feedback received from users, and response to this feedback through program improvements.

Suggested goal: Facilitate AI applications discovery and implementation in the short- and long-term. The NAIRR should track use of resources for AI applications with high social impact -- for example, forecasting climate disaster likelihood with better accuracy and resolution, and precision medicine (i.e. biological datasets to design medical treatment tailored to individuals). It should also monitor use for short- (1-2 years from application), medium- (3-5 years from application) and long-term (6-10 years from application) projects, and evaluate whether there is an optimal ratio of projects falling into these buckets.
Metrics for this goal:

- # of milestones reached or applications discovered related to NAE Grand Challenges
- # of publications from NAIRR use that contribute to projects across diverse fields, and # of citations of these publications
- Time horizon for projects and breakdown across NAIRR portfolio between short-, medium-, and long-term projects.

B. Ownership & Administration

We can envision multiple models of government ownership and administration of the NAIRR, including establishment of a new FFRDC to oversee, a public-private partnership, or ownership by a particular agency with oversight and governance through an interagency body and external input. We encourage the Task Force to explore what governance models have been most effective for past public-private and interagency efforts. Wherever this resource is housed, we urge the NSF, in close coordination with OSTP, to ensure alignment with the multitude of other related NSF activities, such as the HPC Consortium, Open OKN, National Compute Research, AI Institutes, and Cloudbank. Similarly, other agencies with Cloud pilot programs should ensure that the NAIRR does not duplicate these efforts (e.g. the NIH STRIDES program) and instead are able to potentially form a more networked approach.

C. Governance & Oversight

Regardless of which agency or organization manages the NAIRR, we recommend that the National Science and Technology Council oversee its operations and ensure that all relevant agencies are encouraged to contribute data and tools to the resource, that government agencies are focused on using the resource for priority research and application areas, and that NAIRR implementation takes into account lessons learned from other federal cloud programs. We also recommend that the National AI Advisory Council create a subcommittee to provide oversight and feedback to the program, and that members of this subcommittee are chosen to represent the diverse user base for the NAIRR.

The NAIRR should deploy a two-tiered review process to select users for access to NAIRR compute. Most users would apply for "base-level" access to the resource, which would be granted through a light review process. Proposals requiring compute over a certain threshold level would be subjected to a more rigorous merit review process, which would prioritize use cases with the greatest potential to contribute to public knowledge and the public good. *Note that we have not set a threshold for "base level" versus higher levels of access, and would encourage the NAIRR task force to develop a specific recommendation here after consultation with possible applicants.* Other criteria for this merit review should include (but are not limited to): scientific merit; qualifications of the team to undertake the proposed research; whether the proposal is from an underrepresented user group (this could be measured by type of institution, attributes of the individual proposer, geography, field of study); rigor of plans to address AI ethics (privacy, civil rights, fairness, transparency); approaches to minimize carbon impact of the work; originality and other relevant criteria considered for NSF merit review. In addition to criteria for gaining access, the NAIRR should develop criteria for removing users who do not adhere to an AI ethics code of conduct.

Regarding proposal solicitation NAIRR has two primary options: (1) issue and manage its own RFPs, in which case NAIRR will have more control over proposal selection, or (2) serve as a resource for existing research funding agencies to supply their awardees with compute and data. There are advantages to both options and Google does not have a strong preference; *however we do recommend that the NAIRR Task Force determine a best approach early on and articulate this so stakeholders can plan appropriately.*

D. Capabilities required to create and maintain a shared computing infrastructure

The NAIRR should recommend the Resource use open source and open-source-based technologies such as container abstraction layers, open APIs, public codebases, and open source databases. This would ensure operational and technical consistency across public clouds or private data centers and effective management of

infrastructure, applications, and data across the organization. Given that many participating researchers and institutions may be using a particular cloud provider for their research, this would ensure that any applications made available to users are compatible and readily portable in a multi cloud environment further enabling users to move projects across different cloud environments. For researchers that operate on premise, this approach would also allow them to leverage both cloud and on-premise technologies. This strategy enables users to take advantage of all the public cloud services and best-of-breed features according to their needs. Designing an open-source, interoperable platform is critical to achieve this and will also significantly lower the barrier for other users to replicate results (lack of reproducibility has become a bit of a crisis in some fields over the past decade). Google's early adoption of open-source technology has demonstrated that use of open-source and open-source based technologies lead to more innovation, more public benefit and more democratic use of technologies.

Further, the NAIRR implementation plan should adopt a multi-cloud, multi-ML framework-enabled strategy. Multi-cloud is an architectural approach that enables users to leverage the strengths, such as tools and platform services, of multiple providers for various purposes, and gives them freedom from a prescriptive architecture of one single cloud provider. This approach is made possible by using open-source technologies to manage containerized applications. Containers are the important layers for enabling different types of cloud-based software applications. Containerized applications can run on any cloud environment, and even on on-premise technical infrastructures. They separate software applications from the underlying hardware and operating system they run on, allowing them to be deployed in a modular and hardware-agnostic fashion. This would enable interoperability and several important benefits. First, it would facilitate participation of users with different configurations of legacy investments in on-premise technical infrastructure and existing cloud usage. Interoperability would allow users to continue to use their legacy investments to the fullest extent, while taking advantage of future NAIRR cloud-based infrastructure services from a variety of providers. Interoperability also enables more innovation. The less users have to worry about accounting for proprietary idiosyncrasies between the various cloud environments they want to deploy their application with, the more they can focus on their research.

Additionally, we encourage the NAIRR to take a **multi-ML framework approach**: i.e. to consider that there are many cloud AI services (hyperparameter tuning, explainability/interpretability, etc) that do require more interaction with the ML framework than is typically provided by a container. Given that, the roadmap should reflect that the NAIRR should include enabling features to allow TensorFlow, PyTorch, JAX, and other frameworks to be used directly and require containers to adhere to a common specification that would support additional functionality with any ML framework, beyond just training or interference.

Finally, we believe the NAIRR should require that compute resources provided be, at a minimum, carbon neutral (all of the major US cloud contributors have committed to carbon neutrality in their operations; Google has gone a step further and committed to operating on 24/7 carbon-free energy by 2030).

E. Barrier to Accessing Government Data

The NAIRR should facilitate streamlined access for more researchers to both already public and otherwise inaccessible government data (through Data Commons) and to open source software (OSS). Even researchers who are not accessing NAIRR compute should be able to use the NAIRR for data and OSS access (with appropriate privacy protections). Publicly available, non-USG data and OSS should also be included on the platform, including data and open source tools from private sector, international, and state and local government sources, and the NAIRR should serve as an impetus to adopt a common approach to structuring and labeling data (where possible) so it is more useful for researchers (though unstructured data should also be included).

We recommend that the NAIRR co-locate an instance of Data Commons in all NAIRR clouds, which we would provide as an in-kind contribution. This would serve as the vehicle for making existing and newly available public data more useful and for providing a standards-driven data governance process. In order to effectively use the wealth of publicly available data (including data on data.gov) for AI (and other) research, a dataset needs to be processed – this involves locating the data, cleaning it, aligning the schemas of disparate data and ensuring machine readability, etc. This expensive error prone process, which is repeated for each analysis, not only becomes a barrier to the use of data, but also leads to problems of reproducibility in research questions. Cleaning a large dataset is no small feat; before making Google datasets publicly available for the open-source community, we spend hundreds of hours standardizing data and validating quality.

Data Commons does the data processing once and makes the processed data widely available via standard schemas and Cloud APIs. Data Commons is not another repository of data sets (like data.gov or dataverse). Instead, it is a single unified database created by normalizing/aligning the schemas and entity references across these different datasets. So, for example, if a researcher wants the population, violent crime rate and unemployment rate of a county, the researcher does not have to go to three different datasets (Census, FBI and BLS), but can instead, get it from a single database, using one schema, one API. Co-locating updated versions of Data Commons with the NAIRR would therefore enable more effective use of the resource.

The data governance process should also consider whether the dataset is representative of the intended content. Even if data is usable and representative of some situations, it may not be appropriate for every application. Data made available to users should be accompanied by data cards where possible. In addition, NAIRR can ensure users have access to tools like Facets to analyze the makeup of a dataset and evaluate the best ways to put it to use. They should also have access to training on building more representative datasets as well as access to tools such as interfaces like the Crowdsource application. If Google provided Data Commons as a data access platform, Google could also offer some of Google's AI Fairness tools as part of the platform. Similarly, if Google provided access to the Vertex AI platform, NAIRR users could leverage Google Cloud's MLOps tools (including model management), Explanations AI and associated fairness tools, and future model risk management, data and model governance and fairness tools.

The Task Force may also wish to consider using the NAIRR as a central hub for offering restricted access to some types of sensitive government data (for example, health-related data, financial services, and Census data), and to establish a process for granting such access. This could be modeled after the Census Bureau process governing restricted data access.

F. Security requirements & access controls

We recommend the NAIRR support high-level privacy principles and/or risk-based control frameworks that guide the inclusion of cloud service providers (CSPs). NAIRR should ensure CSPs adhere to international standards, such as the need to be certified against a minimum set of internationally-recognized standards that we anticipate will be available by the time this resource is made broadly available: ISO SC42 standards and the NIST AI Risk Management Framework. Voluntary Industry Codes of Conduct are also a helpful proxy for assessing the extent of a CSPs offerings in this area.

The success of a research initiative potentially involving sensitive data depends upon the ability to reliably credential users and provide granular access management. These challenges are further complicated in contexts that require that credentialing and access be managed across a range of research institutions, resource providers, and data sources. To address these complexities and the need to ensure a high level of data privacy and security, **NAIRR should consider zero trust principles and architecture** which provide greater security by requiring parties accessing data to demonstrate that they are who they say they are based on multiple, context-aware signals. **Embracing principles of least privilege**, in which parties accessing data are given access only to the resources that are needed to complete the task can similarly help balance NAIRR's grand research objectives with the need to maintain data and security. **Automated identity and access**

management systems can support the implementation of a strategy based upon both sets of principles at scale.

Finally, we recommend NAIRR consider including technologies that simplify the compliance configuration process and provide seamless platform compatibility between government and commercial cloud environments. The benefit of this is that it can quickly and easily create controlled environments where U.S. data location and personnel access controls are automatically enforced in any of our U.S. cloud regions. For example, Google's Assured Workloads can help NAIRR meet the high security and compliance standards through simple controls that help customers prevent misconfigurations and be more confident in compliance. Assured Workloads can be configured to meet a range of compliance requirements such as those set forth by the Department of Defense (i.e., IL4), the FBI's Criminal Justice Information Services Division (CJIS), and the Federal Risk and Authorization Management Program (FedRAMP).

G. Privacy & Civil Rights

Ideally, the NAIRR will give users access to new AI capabilities, including by expanding access to data such as government data that has been historically difficult to access but that would enable further research into key areas such as bias and fairness. This is because sharing certain types of government data and providing powerful AI tools raises important questions about how to protect people's privacy and rights. There are a number of steps the NAIRR can take to protect privacy and civil rights, including ensuring the appropriate expertise amongst its staff, reviewers, and users, and evaluating proposals for privacy and civil rights protections. These steps are outlined in more detail in our answer to question (3).

Public Cloud providers continue to develop innovative and **up-to-date privacy protections and emerging techniques to learn from sensitive data.** For example, Federated Learning is a technique for training global ML models without data ever leaving a person's device, which has been made available through open-source tools, such as TensorFlow Federated. Another technique is Differential Privacy, which can offer strong guarantees that training data details aren't inappropriately exposed in ML models. Additionally, researchers are experimenting more and more with using **small training datasets and zero-shot learning.**

Similarly, modern public cloud environments provide robust security by design, with native security capabilities that can help cyber defenders address weaknesses and capability gaps in existing security efforts. Too many legacy, on-prem systems continue to be expensive to maintain and hard to secure. Traditional security approaches of continually adding more tools, more people, more compliance have not worked. Furthermore, public

cloud providers can provide innovative solutions like confidential computing, which extends the protections provided by encryption to situations where data is in use.

H. Sustaining the Resource

We believe the NAIRR should be a multi-cloud hosting platform for commercial Cloud resources (as opposed to a new Cloud platform developed by government or academia). Building a new platform from the ground up would require a huge investment of dollars and expertise, and even once built would not have the advantages brought by the scale of existing Cloud providers (e.g. security, operational, and energy efficiency). A multi-cloud hosting platform model would allow the USG to negotiate rates and in-kind contributions from private sector partners. As outlined above, we believe that both academic and private sector users should have access to the same resources -- but not at the same cost. Rates should be lower and subsidized by the USG for academic and government users. Furthermore, as noted above, use should be prioritized based on the public benefit of the proposed research, including plans to publish the work in an open access format (either through an open access publication, journal, or version on arxiv or a similar platform). Any user who does not plan to publish the results of their work should not receive priority, and should be required to pay the full cost of the services.

In order to achieve significant impact, we recommend that the USG fund the resource at $500 million/year or more. In addition, private sector participants could provide in-kind support for the program, for example through low-cost access for certain kinds of users (in particular academics and government researchers), training, and data. As mentioned above, Google would like to offer updated versions of Data Commons as an in-kind contribution. With program funding, the NAIRR should aim to dedicate a relatively small portion of the compute resources (e.g. 30%) to cutting-edge, large-scale ML research (i.e. requiring 1 exaflop+ of compute), and reserve the remaining resources for small- to medium-scale projects. We estimate that this structure would allow for a handful of large projects per year, and thousands of smaller-scale projects.

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

NAIRR should prioritize facilitating public access to and the use of government data sets (including US federal, state, and local government data) through an updated version of Data Commons that is co-located with the NAIRR Cloud(s) to support the design, development, deployment, and operation of AI applications (more details under question 1(E).

In addition to data governance, the NAIRR should prioritize compute access for (a) a diverse group of users and (b) use cases that are likely to lead to public benefit. A merit review process for applications (described in more detail in section 1(c)) should consider representation from federal, local and state governments; a diverse range of academic institutions, and startups. Public benefit should be measured by applicant plans to publish their work, by the likely contribution of the work to addressing social issues, and by the ethics considerations outlined in the proposal (see above).

Finally, NAIRR should provide access for users to existing, open access ML models, which users can fine tune for specific applications. This is important, because it will allow researchers to build on existing models and make more efficient use of compute resources.

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

There are a number of measures that the NAIRR can adopt to reinforce AI ethics and responsibility. For example, the NAIRR should:

- Ensure ongoing engagement with government, civil society, and industry bodies working on AI ethics around the world to share and learn from experiences;
- Ensure in-house expertise spanning relevant areas: technical, ethics, human / civil rights, social scientists, Responsible AI researchers, cloud, legal, and public policy;
- Develop a code of conduct and training for all users, and ensure users demonstrate awareness and training in best practices for privacy and civil rights-related issues;
- Design review process(es) to evaluate individual use cases and build a knowledge base of best practices over time, with defined standing committee members pulling from in-house expertise above, ensuring technical and non-technical members that are multi-disciplinary and include social scientists, human rights experts, and tech ethicists;
- Appoint an advisory committee with representation from widely-recognized Responsible AI experts--both technical and non-technical--across academia, tech ethics, human rights, civil society and industry, and lead staff member to coordinate engagement processes on AI ethics issues;
- Develop and require AI ethics training for government-funded researchers (analogous to research ethics training required for bioscience researchers funded by the NIH);
- Require that proposals include a section on ethics, which can reference institutional ethics review processes, and encourage applicants to subject