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

INTRODUCTION

Booz Allen Hamilton (Booz Allen) is pleased to submit our response to the OSTP/NSF request for information (RFI) for the National AI Research Resource implementation plan (NAIRR). As the largest provider of artificial intelligence services for the Federal government[1], Booz Allen provides professional and technical services to research, design, architect, engineer and integrate AI solutions to accomplish critical missions and maintain U.S. technological leadership. We support some of our Nation's most high profile and innovative programs—including the Joint Artificial Intelligence Center (JAIC), Office of the National Coordinator for Health Information Technology, and the Defense Threat Reduction Agency (DTRA) Operations and Integration Directorate—to transform and advance their enterprise AI initiatives in a deliberate, outcome-focused manner to drive mission impact. More broadly, Booz Allen's AI business encompasses:

- An industry leading portfolio of AI/ML projects across civilian, intelligence and defense organizations ranging from early research to large-scale enterprise operations
- Award winning AI research and development teams publishing in top academic journals and forums[2]
- A Tech Scouting network and unique partnerships with big-tech AI/ML vendors and non-traditional start-ups (e.g., NVIDIA Consulting Partner of the Year 2018-2020; Databricks Federal Partner of Year 2021, AWS ML and MLOps competentices)

Our work recognizes that AI is not a single breakthrough technology, but a complex integration of people, processes, and technologies with a responsibility to use AI in a way that centers around people.

1.0 WHAT OPTIONS SHOULD THE TASK FORCE CONSIDER FOR ANY ROADMAP ELEMENTS A THROUGH I ABOVE, AND WHY?

1.1 {Implementation Roadmap-A} Goals for establishment and sustainment of a NAIRR and metrics for success

Booz Allen suggests successful establishment and sustainment will revolve around three core attributes:

- **SHARED INFRASTRUCTURE**: Create, implement, and manage compute infrastructure resources that serve as a common sandbox for broader AI R&D and is representative of the changing operational environment with measurements to assess speed, progress, and sharing.
- **TOOLS, METHODS & DATA:** Make available a wide variety of rich data sets, world-class algorithms, and example uses cases that, when combined with the computing resources, will better democratize AI and allow for more contextual understanding and development. Measurements should include the amount and type of data sets (to include synthetic) and the amounts and types of openly available models, as well as the use/consumption of both the data and models.
- **EDUCATION & TRAINING:** Ensure AI-related assets are widely available to existing users while accommodating a potentially rapid growth rate; etc. This should include prioritizing education, training, and opportunities to engage repeatedly with real mission users to discuss the outcomes needed to achieve adoption. Measures should include the amounts and types of engagement and learning types.

[1] Bloomberg Government Market Analysis

[2] Thought Leadership for AI, Analytics, and Data Science (boozallen.com)

1.2 {Implementation Roadmap-C} A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources

We propose the creation of the NAIRR Council to coordinate between the organizations that hold a stake in the success of National AI research initiatives. This Council could follow the mold of other high-level coordinating organizations such as the National Security Council (NSC) and Joint AI Center (JAIC), and would facilitate working groups, establish strategic priorities, coordinate resources, and direct stakeholder efforts. This council can provide oversight to ensure progress against strategic goals and clear priorities. In this unprecedented partnership across industry, academia, and government, the council should be proactive, outcome-focused, and sensitive to feedback. They would allocate resources judiciously by providing and maintaining clear priorities. This will include leveraging Agile principles for adoption to consistently prioritize, groom, and refine the roadmap with end user feedback. As a part of the "Council" charter, they could define clear roles across decision-makers to best leverage strengths. Additionally, they could develop governance, IP protection policies and invest in researchers through the creation of public-private partnerships that enable technological innovation and collaboration by deploying best-practices and protecting data and technology[3].

To build strong and collaborative public-private partnerships, we recommend: (1) **Accountability**: as partnerships evolve and collaborate, it's important to have policies in place that hold the partnership accountable for successful collaboration[4]; (2) **Diversity**: collaborating with diverse companies and academic partners is a catalyst for innovation and rapid progress; and (3) **Integration**: provide a marketplace for AI and other strategic technologies to be implemented within the government through strategic investments in public-private partnerships to drive R&D[5].

Recommended key functions: (1) Establish and review metrics and measurements of success; (2) Establish and review roadmap/implementation plan to ensure progress; (3) Continuously prioritize and update "backlog" based on feedback, user-testing, progress, and data; (4) Troubleshoot any problems that may arise hindering progress; and (5) Report to Congress and brief external stakeholders

1.3 {Implementation Roadmap-D} Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country…educational tools and services…and scalability of such infrastructure

In our recently published O'Reilly report, Enterprise AI Operations, we highlight the capabilities and components needed to successfully adapt and employ enterprise-wide artificial intelligence capabilities.[6] A mature AI operations pipeline that will ensure scalable infrastructure and enable analytics and AI for NAIRR research needs should be thought through from the beginning to avoid rework and better support outcomes.

For more specific capability descriptions, see our response to Question 6. Below are some key overall areas associated with creating a framework to take capabilities from the lab to operations for NAIRR consideration.

- **Responsible AI Adoption** | Ensure AI solutions, when deployed, meet performance requirements, adhere to organizational standards and values, and are designed for adoption to achieve real mission outcomes.
- **AIOps**| Integrate processes, strategies, and frameworks to operationalize AI and address real-world challenges and realize high-impact, enduring outcomes.

---

[3] National Security Commission on AI Final Report (NSCAI), page 205

[4] Casady, Carter, et.al., "A 'New Governance' Approach to Public-Private Partnerships: Lessons for the Public Sector," 2017, Stanford University."

[5] National Security Commission on AI Final Report (NSCAI), pg. 449

[6] Booz Allen Hamilton, "Enterprise AIOps: A Framework for Enabling Artificial Intelligence," O'Reilly Report, August 2021.

- **Data Engineering and Data Operations (DataOps)** | Locate required data and develop repeatable pipelines to increase value and make enterprise data accessible while promoting re-use.
- **ML Engineering and ML Operations (MLOps)** | Develop advanced algorithms using supervised, unsupervised, reinforcement, deep learning, etc. as required to support complex decision making.
- **Systems Engineering and DevSecOps** | Apply a structured framework for integration, documentation, and automation to develop, deploy, and monitor software and system solutions across an organization. Development, security, and operations (DevSecOps) integrate the critical components of security and focus on operationalizing applications through software and systems engineering.
- **Reliability Engineering** | Establish focused, clear objectives for AI solutions that are realistic with well-defined and quantifiable measures of success.
- **Infrastructure and Cybersecurity Engineering** | Protect data and AI applications for the long-term success of operationalizing AI with a strong technical architecture and cybersecurity policies.

1.4 {Implementation Roadmap-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

See our response to Question 6 below. We recommend looking at data, compute environment, and AI development holistically rather than independently.

2.0 WHICH CAPABILITIES AND SERVICES PROVIDED THROUGH THE NAIRR SHOULD BE PRIORITIZED?

One challenging aspect when establishing the NAIRR is appropriately bounding services. Put simply, NAIRR shouldn't try to become all things to all people. To establish and implement NAIRR at a rapid pace, we recommend focusing on three key areas: standardizing an AI Reference Architecture needed for a compute infrastructure, providing data sets and compute resources, and making available educational tools and resources.

2.1 AIOps—AI Reference Architecture for a Compute Infrastructure

To operationalize and deliver responsible, scalable AI solutions, Booz Allen developed an approach that starts with our recently released AI Reference Architecture (RA). A reference architecture provides technology agnostic guiding principles to standardize and accelerate the delivery of AI through common components, capabilities, processes, and terminology (Figure 1). Adopting a standard RA enables an organization to produce real-world solutions from an abstract framework, which drives consistency and standardization, surfacing risk early and supporting mitigation and reduction measures.

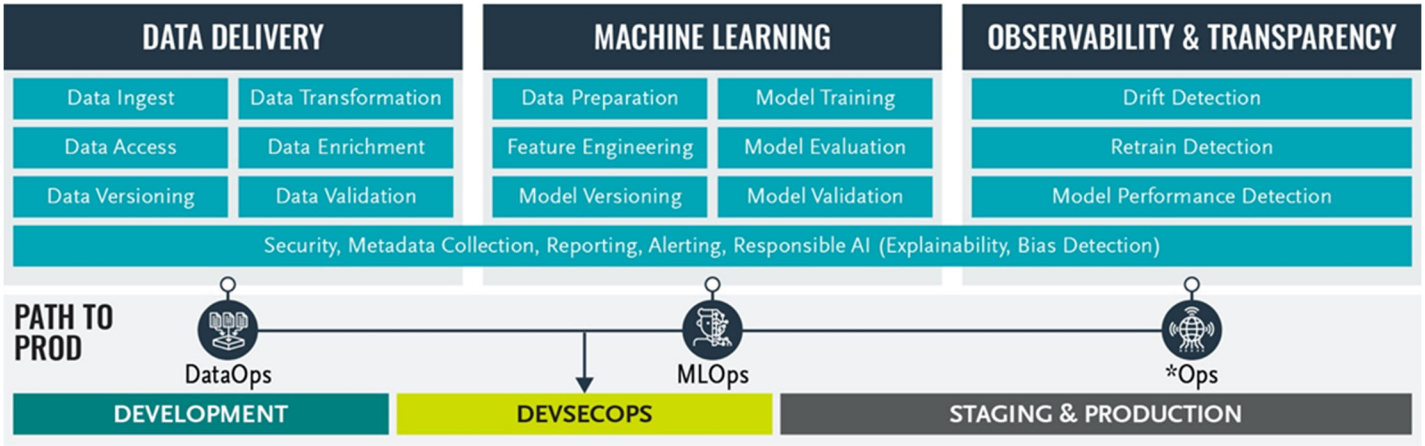

FIGURE 1: BOOZ ALLEN'S REFERENCE ARCHITECTURE (C) 2021 BOOZ ALLEN HAMILTON. ALL RIGHTS RESERVED.

- **DATA DELIVERY**: Covers activities that span ingestion, storage, transformation, enrichment, and delivery of data for analytics at scale. Robust, repeatable data delivery provides a critical foundation for analytics activities, and provides lineage and provenance collection to support data governance.
- **MACHINE LEARNING**: Encompasses the performed activities to prepare and transform data into insights and inferences that serve client and other business needs. The RA defines the main capabilities that when combined comprise a holistic ML Workflow.
- **OBSERVABILITY AND TRACEABILITY**: Provides processes and capabilities for monitoring and reacting to changes in ML workflow performance. This includes tracking performance, detecting when performance shifts beyond acceptable limits, and triggering appropriate actions in response.
- **CROSSCUTTING COMPONENTS**: Provides core functions including security, metadata collection and explainability to support the development of analytic services.
- **PATH TO PRODUCTION**: Combines software delivery best practices with Data Operations (DataOps), Machine Learning Operations (MLOps) to establish a robust path to production.

2.2 GOVERNMENT CURATED DATA SETS AND COMPUTE RESOURCES {ITEM D}

Data and computational resources are foundational to successfully creating responsible AI; therefore, related capabilities and services should have priority. This should include data, curation, hosting, auto tagging, maintenance of a variety of types of data sets for different types of use cases that have been cyber secured. Additionally, we recommend incentivizing the sharing of data sets for research purposes. We would also suggest collaboration with the Office of the National Coordination for Health Information Technology (ONC) on their recent report "Training Data for Machine Learning (ML) to Enhance Patient-Centered Outcomes Research (PCOR) Data Infrastructure"[7] as the main goal for this effort is to make data sets available for AI/ML research. Additional information is in our response to Question 1.4 Implementation Roadmap D above and Question 6.

2.3 EDUCATIONAL TOOLS AND RESOURCES {Item D}

Getting to adoption is more than just providing data and computing resources. It's also creating a culture of learning and sharing, which includes datasets, models and other resources that emphasize an open-source mentality. Prioritizing retraining and collaboration with educational institutions will help integrate AI curriculums into everyday learning, starting from a young age with reinforcement as educational groups move through their learning journey. It's crucial to remove the barrier to entry for learning about and engaging with AI to shape the workforce required for the United States to stay at the forefront of AI innovation. For example, Booz Allen has invested in "The AI Education Project" to help lower barriers to AI education. The project works with K-12 schools to increase access to computer science education to help students build a foundational knowledge of how AI works and how it affects their lives[8]. Introducing AI early in a student's education helps remove AI's mystique and further integrates it into society as a fundamental tool. Partnering with a range of higher-education institutions will provide high-quality technology education to foster a creative, innovative, and AI-fluent workforce.[9] The University of Florida entered a public-private partnership with NVIDIA to build an AI-centric data center where students receive hands on experience with AI and industry leading tools like

[7] The Office of the National Coordinator for Health Information Technology-"Training Data for Machine Learning (ML) to Enhance Patient-Centered Outcomes Research (PCOR) Data Infrastructure" Report (healthit.gov)
[8] https://medium.com/the-ai-education-project/introducing-the-ai-education-project-3c1f1fc31fd2
[9] National Security Commission on AI Final Report (NSCAI), p. 543

supercomputers[10], advancing research efforts. Exposure to environments like a large, shared infrastructure, AI research programs, and curriculum will help solidify U.S. colleges as technology innovation hotbeds**.**

3.0 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?

Responsible AI involves more than just fairness and ethical outcomes. These concepts are important, but they only represent a system's characteristics—which are neither good nor bad on their own. As a result, assessing whether an AI system meets responsible and ethical requirements requires two key elements: 1) clearly articulating what responsible AI development looks like and adopting from inception, and 2) incorporating risk mitigation strategies to identify, assess, and address possible bias.

3.1 CLEARLY ARTICULATE RESPONSIBLE AI DEVELOPMENT

Thinking about responsible AI from the design phase all the way through to implementation and monitoring is critical to ensure it represents ethical outcomes.[11] Responsible development is a core aspect that needs to be at the forefront—focusing on the combination of elements that make up responsible AI, such as adoption, ethics, and the workforce.

Responsible AI becomes meaningful only when considering how an artificial intelligence system impacts people. The "*who"* question is critical. In other words, responsible AI needs a focus—whether that's an individual, a group or class of people, or an entire society. But responsible AI *also* needs a source of authority, or moral compass. That can be an organization's shared values and principles, or a society's norms and cultural practices. It is only then that we can address, for example, what a "*fair*" AI system is, for whom, and why we consider the implications to be fair in the first place. Embedding these values into the governance structure with the proper controls and measures to validate execution is key. This governance structure, including controls and metrics, should have agreement at an organizational level and not on a per project basis.

There are five main considerations when integrating Responsible AI:

1. Understand how artificial intelligence systems will impact an organization's stakeholders in specific and tangible ways. This assessment should include routinely considering the political, economic, social, technological, legal, and environmental impacts across different stakeholders over time.
2. Build meaningfully diverse and inclusive development teams. Include members with different backgrounds, skills, and thinking styles. A team's collective experience and insights will reduce unconscious bias, identify potential unintended consequences, and better reflect stakeholders' wide-ranging values and concerns.
3. Develop mechanisms for data provenance and auditability to verify AI systems are operating as intended. If something goes wrong, data tracing and auditability mechanisms will help uncover data or concept drift or potentially expose upstream/downstream data issues. Clear accountability mechanisms and data test can help reduce ethical concerns (e.g., data bias and amplification of the bias in ML for training, inference, etc.), so it is critical to transparently account for the results. Teams should understand that they are accountable for the actions, outputs, and impact of their models.

[10] University of Florida, UF Announces $70 million artificial intelligence partnership with NVIDIA, v July 2020.
[11] Booz Allen Hamilton, "Enterprise AIOps: A Framework for Enabling Artificial Intelligence," O'Reilly Report, August 2021

4. Stay informed regarding AI technical developments. Because this field of study changes rapidly, tools used to design and implement ethical systems have limited shelf-lives. A model's sophistication will often outpace ethical tools, increasing the probability something will go wrong and reducing the ability to fix it if it does. Maintaining awareness of AI technical developments and implementing measures to monitor can help mitigate risk and protect an organization from unintended consequences.
5. Design systems with specific applications and use cases in mind. Assessing the "fairness" of a model requires *context* and *specificity*. AI systems should be fair, but fair to *whom*? And in *what way*? Fairness is a laudable goal but only becomes useful when applied to a specific situation. Something that may be a fair outcome for someone in one situation could appear totally unfair in another situation.

3.2 INCORPORATE RISK MITIGATION STRATEGIES TO IDENTIFY, ASSESS, AND REDRESS POSSIBLE BIAS

As described in our recent "NIST Artificial Intelligence Risk Management Framework Request for Information (86 FR 40810)" response, we are deeply committed to maintaining an acute awareness of the societal and environmental impacts of our AI systems and applications, and we ensure that our company and our people design AI systems that are grounded in real-world implications. To provide checks and balances and ensure the implementation of AI guiding principles, we believe that the implementation of an AI governance process, like that described in Figure 2, is necessary. In addition to ensuring the evaluation of AI projects to systematically mitigate risk, governance builds stakeholder confidence in AI through an organizations' responsible use of AI. Well-designed controls promote the application of AI through effective management, ensuring that it is meeting performance requirements and ethically used.

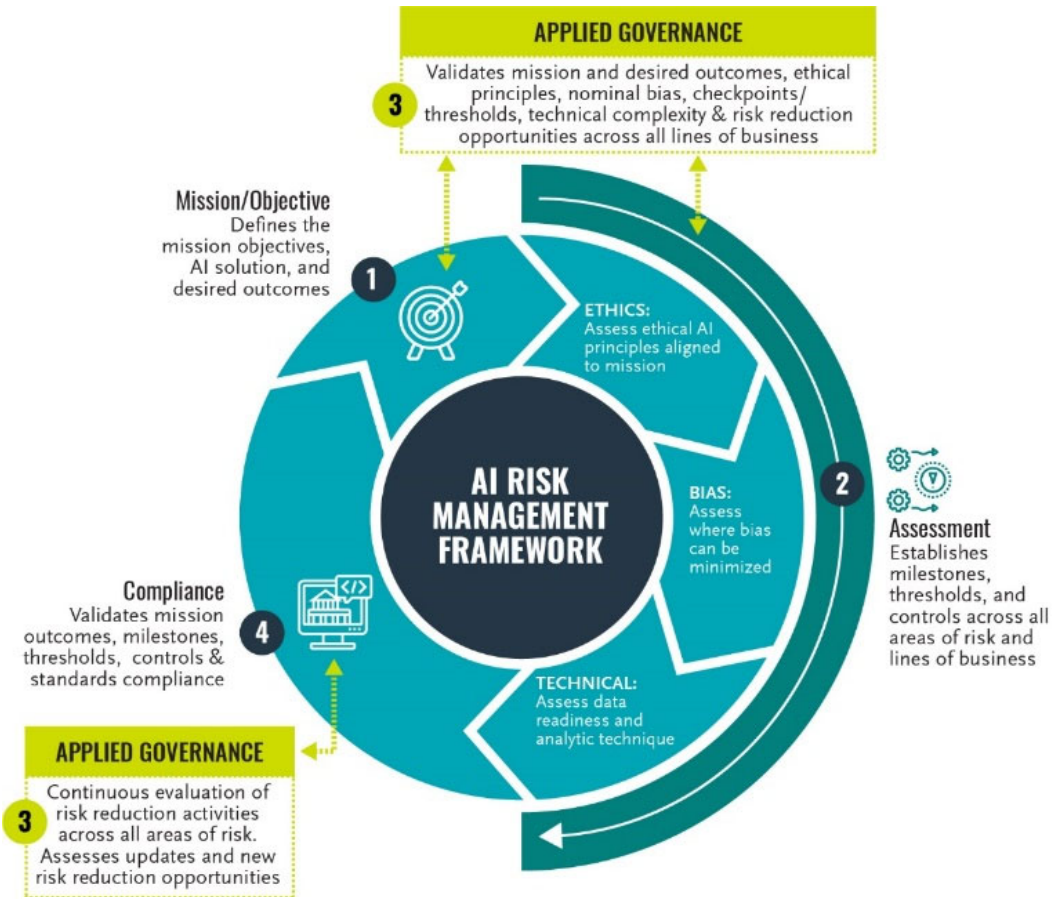

FIGURE 2: BOOZ ALLEN AI RISK MANAGEMENT FRAMEWORK (C) 2021 BOOZ ALLEN HAMILTON. ALL RIGHTS RESERVED"

4.0 WHAT BUILDING BLOCKS ALREADY EXIST FOR THE NAIRR, IN TERMS OF GOVERNMENT, ACADEMIC, OR PRIVATE-SECTOR ACTIVITIES, RESOURCES, AND SERVICES?

The private sector is rapidly pushing for AI growth through forming new partnerships between technology powerhouses and investing in smaller firms and start-ups, creating scalable multi cloud environments,

developing a range of AI frameworks for rapidly developing modeling tools, and creating/implementing large scale workforce programs. Booz Allen is committed to investing in AI innovation, including investing in non-traditional businesses with potential to positively disrupt public sector missions. For example, we invested in Latent AI with the goal to leverage the partnership to help clients implement AI models into sought-after end user devices to help drive AI adoption.[12] Investments and collaboration between private-sector entities foster rapid innovation and implementation of technology at a greater scale.

Workforce training will be one of NAIRR's most powerful tools. Many private-sector companies have established and well-tested training programs to educate, train, and grow their workforce. Through a partnership with NIVIDA's Deep Learning Institute (DLI), Booz Allen has developed a Deep Learning Training service to help clients build data science skills required, to become experts that can implement Machine Learning technology within their company or agency.[13] This type of customized facilitated training service will help accelerate analysis and drive mission success.

There are numerous building blocks that already exist across the Federal Government, Academia, and industry. We would suggest the Government conduct a larger data call to get more information about those efforts. To be most effective, we suggest setting up some overall categories/bounding boxes with clear definitions so that the collective group can better organize their inputs in a digestible manner. New efforts arise every day and would encourage non-traditional thinking of how-to best leverage and learn from those efforts. With that being said, we offer a few specific building blocks below for consideration in addition to the other efforts listed throughout our response:

NAME	BRIEF DESCRIPTION
DoD ADVANA Data Platform	A central hub for advanced analytics, data science and AI connecting Senior Leadership to authoritative data sources needed for data-driven decisions.[14]
VA/DOE- VA's Million Veteran Program	The Department of Veterans Affairs (VA) and the Department of Energy (DOE) are partnering to drive technology innovation and transform health care delivery for Veterans. The partnership brings together VA's healthcare and genomic data with DOE's high-performance computing (HPC), artificial intelligence and data analytics.[15]
The AI Education Project	The AI Education Project is a 501(c)(3) non-profit centering equity and accessibility in AI education. They educate students, especially those disproportionately impacted by AI and automation, with the conceptual knowledge and skills they need to thrive as future workers, creators, consumers, and citizens.[16]
DoD GameChanger Open Sourcing	Over 15 thousand documents govern how the Department of Defense (DoD) operates. These documents exist in different repositories, often on different networks, are discoverable to different communities, updated independently, and evolve rapidly. GAMECHANGER offers a scalable solution with an authoritative corpus comprising a single trusted repository of all statutory and policy driven requirements based on Artificial-Intelligence (AI) enabled technologies.[17]

[12] Booz Allen Hamilton, "Investment builds on AI capabilities supporting algorithmic warfare at the edge," July 2021.
[13] Booz Allen Hamilton, "Deep Learning Training"
[14] Be Ready for What's Next - Government Technology Insider
[15] DOE and VA Team Up to Improve Healthcare for Veterans | Department of Energy
[16] The AI Education Project
[17] GitHub - dod-advana/gamechanger: GAMECHANGER aspires to be the Department's trusted solution for evidence-based, data-driven decision-making across the universe of DoD requirements

Data.gov	The U.S. Government's open data site—access to data, tools, and resources to conduct research, develop web and mobile applications, and design data visualizations.[18]
NIH NLM Data Science Training	NIH National Libraries of Medicine (NLM) Data Science Training Program that provided targeted training to all NLM's 1,700 staff members. With a focus on becoming more aware of data science and its incorporation into so many NLM products and services.[19]

6.0 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?

The quality of resources and investments will have the greatest impact on NAIRR's goal to democratize access to AI R&D. A holistic approach that better connects research to operations will allow for integrating AI R&D into the fabric of America. We encourage the Government to think of both basic and applied research and how to translate and integrate into the operational environment that exists today including the use of "operational test beds." Below are the key areas that we see as critical for democratization, and some (not exhaustive) associated sub areas for consideration.

DATA QUALITY & AVAILABILITY | There is a need for broad and consistent access to high-quality operationally relevant catalogued data for training and testing of AI models. This should include assessing the ability to access data required, manage its quantity, and evaluate its quality, as well as, evaluating the capacity to understand the data (e.g.) data dictionary. Additionally, there is need for the ability to correctly evaluate biased data and make the proper corrections.

- **Data Readiness & Management**: Managing data as an asset can help NAIRR address these challenges and support AI needs through a multilayered approach to data readiness. A holistic approach includes the following key areas: (1) Data strategy: aligns data capabilities with AI needs and establishes effective data management and data usage practices tailored to AI needs, (2) DataOps processes optimize workflows for automated data ingest, processing, and storage to help keep up with data needs at scale, (3) Standardized methods enable data ingest, transformation, validation, and preparation, (4) Data governance expedites value and spans compliance, risk, and regulation related to data (including privacy, security, and access controls), and (5) Responsible data policies specify access rights, "right to see" authorizations, ethical principles, and acceptable applications for data usage across the organization.

- **Synthetic Data**: There is not enough diverse, high-quality, labelled data accessible, thus limiting the application of trained ML models to many domains. Synthetic data is the generation of artificial data with the aim of reproducing the statistical properties of an original dataset generated by real world events. This data can be either partially or fully synthetic, with the former containing generated data alongside original data, and the latter composed exclusively of generated data. It is a powerful solution for those who do not have the resources to collect, label and process huge amounts of data typically needed to train complex algorithms. The Department of Veteran Affairs has been actively using synthetic data through its Veterans

[18] Data.gov

[19] NIH- Building Data Science Expertise at NLM – NLM Musings from the Mezzanine (nih.gov)

Health Administration Innovation Ecosystem (VHA IE) because of restrictions associated with PHI and PII.[20]

- **Open Source:** Maximize data access by streamlining policies and regulations to promote and increase data sharing across organizations. Supporting and investing in AI tools like JAIC's GameChanger[21], which is modernizing data access rules, will ensure more users have access to relevant and high-quality data. Other evolving data sharing practices include multi environment designs based on security where verified researchers could have access to more sensitive information.[22] Incentivizing the private sector and academia to share data sets is a critical element for democratizing AI R&D.

GOVERNANCE| There is a need to ensure the creation/use of "Responsible AI" and risk management frameworks that include oversight and sustainment for all data, model, and computing environment efforts focusing on reusability. This includes processes for configuration management, testing, and verification & validation.

- **Responsible AI- Risk Management**: See Question 3 above and extensive response to the NIST Artificial Intelligence Risk Management Framework Request for Information (86 FR 40810)"[23]
- **Data Test & Evaluation**: Integrated data tagging and labeling capabilities while also providing the ability to develop ground-truth datasets.
- **Model Management**: As with any other form of software, ML models need documentation, testing, and version-control to help the team maintain knowledge of its structure and functionality. Unlike conventional analytics, we must rely on surrounding documentation, testing, and metadata.

ARCHITECTURE & DESIGN| There is a need for a modular software architecture that deployable in a timely manner while ensuring a focus on security and scalability. It needs to include the use of open, secure software architectures comprised of best-in-class COTs, GOTs, and open-source components that balance quick experimentation with affordable scalability. The system should support interoperability and be reliability operated and maintained.

- **Evolutionary data architecture**: Data architectures designed with flexibility and adaptability in mind to evolve at the rapid pace of innovation. Maintaining the proper design abstraction allows for the component replaceability necessary to keep up with AI's rapid evolution.
- **Scalable and Flexible data pipelines architecture**: Designing data pipelines for scale and flexibility at the beginning for efficiency rather than attempting to scale at subsequent stages of growth.
- **Institute open API interfaces for data sharing**: An application programming interface (API) defines how to push or pull data during a data exchange. APIs help bridge data silos and break them down into more usable parts that greatly increase integration and reduce development time.

SECURITY | There is a need for broad and consistent access to configurable, scalable, accredited computing environment to support development, experimentation, and operational hosting of AI capabilities. This requires ensuring systems built are robust and secure and incorporate security practices at every stage of development. Additionally, there should be methods to properly monitor new complex AI solutions and have techniques in place to combat adversarial AI.

---

[20] How synthetic data will improve Veteran health and care - VAntage Point
[21] National Security Commission on AI Final Report (NSCAI) pg. 305
[22] National Security Commission on AI Final Report (NSCAI) pg. 449
[23] https://www.regulations.gov/comment/NIST-2021-0004-0058

- **Adversarial-** As more and more systems integrate AI, there is growing concern about how to protect these systems from attacks (e.g., evasion, poisoning, model-stealing, backdoor, etc.). This requires a novel approach to AI protections using methods from the well-established domain of control theory to provide mathematically provable protections built into the system during development. This provides engineers with a flexible and generalizable way to design multi-layer neural networks that are resilient against many attacks without compromising the system's ability to learn.

- **Securing the Infrastructure**: Application and mission owners must architect their cloud environments to not only be compliant with security and risk standards but simultaneously ensure the implementation of the right cybersecurity designs and tools to protect against emerging threats and attacks. This requires continued enhancements as the platform becomes operationalized, resulting in the need for organizations to migrate to a SecDevOps continuous development culture**.**

CULTURAL & ADOPTION| There is a need for a clearly articulated AI vision and strategy that directly supports the organization's mission. This should include proper training, communication plans to support building of trust and transparency, and operational user involvement and feedback.

- **Education & Training**: There is not enough AI-fluent talent to meet the government's needs and a rapidly evolving technology landscape. Workforce training and educational tools are key to supporting a strong computing infrastructure. A focus on training so that all users are AI knowledgeable also provides access to a larger talent pool, to build meaningfully diverse and inclusive design and development teams. This should include educational and training courses ranging from intro to refresh course at varying degrees of technical competency based on roles individuals will play focusing heavily on achievable mission outcomes.

- **Mission Acceleration & Mission First Engineering**: To truly bridge the "compute divide," the focus needs to also include the connection to the mission so that the collective community (research and mission) can learn and grow together. Many are too focused on protypes and pilots that when brought to operations may never actually deploy as the operational environment was not part of the original design. The creation of AI is most effective when starting with the outcomes desired and fully understanding the "mission" needs. An AI team should be a cross-functional, integrated team working on holistic AI/ML solutions that leverage the expertise of the members (technical, operational, enabling) with the key understanding that all functional areas are critical for success. Operational feedback loop(s) and learning is one of the most essential pieces of getting AI into operations. Your teams must be able to monitor data, models, applications, and processes to evaluate potential changes/updates needed throughout their life cycle to ensure you are achieving the right outcomes responsibly. Like sports, we need to practice like we will play so connections to operational environments from both a testbed environment and connection to operational users is critical. Keeping research separate will continue to exacerbate the divide between getting to operational AI and real adoption with mission outcomes.