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

Palantir

EXECUTIVE SUMMARY

Palantir Technologies Inc. ("Palantir") is a software company that provides data integration, analysis, and decision-making platforms. Our platforms are used as a data foundation and development infrastructure for artificial intelligence (AI), giving us unique insight into the challenges and successes of AI programs. Our software supports AI research on topics ranging from pharmaceutical drug combinations to COVID research.

We applaud the National Artificial Intelligence Research Resource (NAIRR) Task Force ("Task Force") for identifying and advocating for a holistic computing ecosystem to support AI researchers, practitioners, and students via the implementation of a National Artificial Intelligence Research Resource. In our experience, data – the quantity, quality, and appropriateness to the problem – is the greatest determinant of an AI effort's outcomes and should be central to the NAIRR's ultimate implementation.

To that end, we encourage the NAIRR to focus on ensuring the completeness, integrity, trustworthiness, security, and proper use of data used for AI research by prioritizing infrastructure, governance, and intentional resource allocation.

Infrastructure and Data Foundation. Trustworthy and secure AI models require a trustworthy and secure data foundation (as the saying goes, "junk in, junk out"). In considering the NAIRR implementation roadmap, we urge the Task Force to prioritize specific infrastructure components that are most essential and practical to ensuring a quality AI output: intuitive data integration and pipeline transformations, audit logs, granular access controls, the ability to version and branch data, and collaboration features for annotating datasets and identifying addressable issues over time (e.g., statistical and other forms of unwanted data bias). With these key investments, the NAIRR effort has the potential to support, enrich, and grow the AI research community by providing a shared data environment, not just static datasets. The NAIRR data resource can enable tangible, trustworthy results through collaboration tools that allow for secure modifications, annotations, and improvements to datasets and the ability to share knowledge and performance metrics. These investments will enable the NAIRR infrastructure to scale and support cross-organization, cross-discipline research to springboard the U.S.' AI capabilities.

Technical, Governance, and Cultural Awareness in Responsible AI. Not only should the NAIRR incorporate best practices around AI Privacy and Civil Liberties (PCL), Bias, and Ethics, but it can — and should — enforce these best practices (e.g., stringent security and access controls) through a combination of a) technical infrastructure that facilitates enforcement of key data protection and responsible AI principles; b) governance that incentivizes and rewards best practices; and c) cultural awareness and discipline-specific frameworks to contextualize AI research and guide determinations of when and how AI applications can best align with the interests of impacted communities. Through our decades-long work developing and implementing data integration and management platforms built on granular security and privacy-preserving capabilities, we have demonstrated that well-engineered data infrastructure must provide the technical implementation measures that attach to and reinforce critical institutional governance measures. The challenges that AI research will ultimately address are not exclusively

technological in nature, rather, they are techno-social and require an understanding of the cultural contexts, innovative data science practices, and institutional controls.

Problem Prioritization through Resource Allocation. The Task Force is well-positioned to provide oversight driving NAIRR use cases in an opinionated way towards pressing, relevant problems and historically overlooked or underfunded initiatives. Similar to how the National Science Foundation approves grants, the Task Force can provide incentives and allocate resources (compute, data, etc.) based on problem evaluation. For example, the Task Force could identify AI safety research as a topic requiring more focused research and development. When industry alone researches this topic, it is likely to direct its attention to narrow, commercially-focused use cases. The NAIRR, in contrast, could expand the scope and insights produced by the AI research community. With the appropriate infrastructure and governance in place, the NAIRR has the potential to shed light on how AI resources are being used, discover whether the most important problems (as seen through a broader societal impact lens) are being researched, drive investment, and compound knowledge from across academic, government, and industry participants.

Categorizing AI Research & Development
The themes above can manifest differently based on the category of AI research, development, and application in question. We recommend that the Task Force clarify which of these categories NAIRR aims to address, recognizing each has different data, infrastructure, and organizational needs.
1) ***Pure or Basic AI research and development*** focuses on advancing the state of the art of techniques and methodologies for artificial intelligence. Basic research requires access to compute, ability to build new network types, and scaling infrastructure for large datasets.
2) ***Applied AI research*** focuses on taking an existing AI algorithmic approach and exploring its utility in the context of a real-world problem. Applied research requires the ability to collaborate with non-coding subject matter experts (biologists, physicists, etc.) as well as access standard libraries (pytorch, tensorflow).
3) ***Operational AI*** productionizes use of AI using real-world data, with real-world outcomes. Operational AI requires DevOps tooling to deploy, scale, and monitor models.

OUR PROPOSED VISION FOR A NAIRR DATA INFRASTRUCTURE

We encourage the Task Force to commit to investing in the three areas mentioned above (infrastructure, governance, and problem prioritization) as part of the NAIRR implementation roadmap to give researchers and students access to a powerful, trustworthy, and secure end-to-end AI research environment. Elements of this vision are already a reality at NIH and other agencies (see our response to Question 5 on page 9), and the NAIRR should evaluate the feasibility and sustainability of incorporating these practices.

Ability to search for, select, and combine data: As stated in the introduction, data and the associated data infrastructure are the greatest determinants of an AI program's success or failure. The NAIRR should empower AI researchers to:

- **Easily discover potential training datasets relevant to their use case** through a transparent data catalog, as well as metrics and metadata on those datasets.
- **Branch a data set**, modify it, and make those modifications available to the broader research

community in a fully secure and transparent way. Branching datasets supports discovering, recording, and mitigating bias and other data issues for the entire research community.

- **Collaborate** on training data sets without compromising data lineage, integrity, or security.
- **Granularly secure data** with low-friction, built-in access control tooling. Not all datasets should be fully accessible to all users (e.g., to protect sensitive information) and oversight is required to assess the potential consequences of combining data sets.

Access to the platform through a straightforward registration process. New users should ideally be able to access and begin using tools the day that they sign up. A complex and/or drawn-out registration process will likely alienate some of the communities that the NAIRR is most seeking to serve. Automated security validations that are built into the technical infrastructure and a user-friendly interface will be key to enabling this step.

Ability to train, test, and retrain AI models. The NAIRR should provide researchers with the development environment, computational power, and tools required to train AI models in a streamlined manner, or enable them to use their own tools, connected via open APIs. Users should be able to easily capture and share their work products (libraries, models, data modifications) back to the NAIRR, enabling knowledge to compound over time. A technical infrastructure designed to capture and share knowledge enables appropriate contextualization of existing data and research, which in turn enables directing resources to priority problem areas.

Ability to evaluate AI model performance within the NAIRR platform. Researchers should be able to view and capture metrics that show the performance of their AI models as well as how a model compares to baseline standards, and to capture performance against different evaluation datasets. Researchers could have the option to make this performance data discoverable so that others can learn from their work, while also securing their data so that it is only accessible by those with appropriate permissions.

Ability to contribute back to the NAIRR. User friendly tools should be provided that incentivize researchers to contribute back to the broader NAIRR community. For example, researchers could make available new and improved data pipelines, data annotations, model templates, analysis frameworks, etc. Incentives such as additional compute allocations that encourage contributing knowledge back to the AI research community would allow the collective knowledge of the NAIRR compound over time and streamline future research efforts.

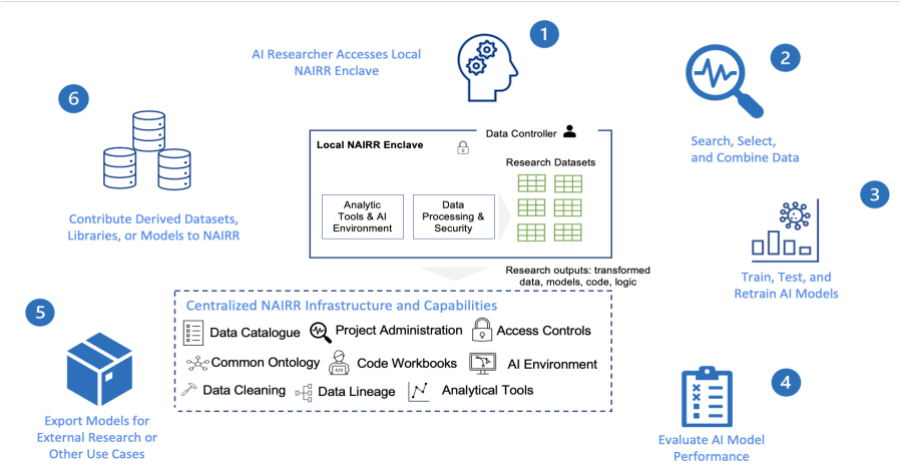

Figure 1: This graphic depicts the AI research lifecycle as enabled by a NAIRR digital infrastructure.

RESPONSES TO QUESTIONS

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

In general, the "north star" for the NAIRR should be implementation of and provision of access to a data infrastructure that facilitates and democratizes AI research. Reflecting on our experience in enterprise data management, Palantir suggests the government consider the following NAIRR implementation goals:

Goal	Benchmarks
To provide the U.S. research community with high-quality, broadly available data sets that have been screened for common forms of statistical bias for AI model training and evaluation.	• Researchers have access to quantitative tools and techniques for assessing the quality of datasets and their effect on model outputs. • The NAIRR continuously integrates advancements in data evaluation tools and techniques.
To accelerate research of AI techniques and models that address critically important problem sets in both the public and private sectors.	• NAIRR infrastructure enables collection, reporting, and surfacing of metrics showing where research is being conducted. • NAIRR program office has access to these metrics and can ID research gaps.
To measure, and demonstrably improve, the effectiveness, accuracy, and fairness of AI models over time.	• The NAIRR platform contains infrastructure to calculate and evaluate the performance of AI models. • Metrics for AI effectiveness and accuracy are transparently and accountably recorded to curate the specific research area, along with other critical, contextual, and social relevant evaluations (e.g., fairness metrics)
To expand the AI research community to include researchers from historically disadvantaged communities.	• Metrics about who is using the NAIRR and what kind of research is being conducted are collected in the NAIRR platform (within the bounds of a consent framework and protected by access controls). Over time, metrics show increased participation from historically underserved communities.

1.D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country.

To create and maintain a shared AI research computing infrastructure that delivers access to quality and trustworthy data, we recommend that the NAIRR contain the following capabilities:

A dynamic, flexible, and secure data environment. Without integrated, clean data, model outputs are likely to be incomplete, inaccurate, and to perpetuate system flaws (including forms of bias). The NAIRR must have the capability for users of varying technical backgrounds to integrate any type of data source and require including critical information on data quality, provenance, and data pipeline health to protect against faulty input data.

Privacy and security controls. The Task Force should strive for maximum granularity in the NAIRR access controls, down to the row and column levels when required, to ensure that data will not be misused while also enabling the greatest flexibility in collaboration. Additionally, transparent propagation of security controls to downstream datasets will allow researchers to appropriately discover data used to build pipelines as required to pursue a properly scoped AI project. The NAIRR digital infrastructure should also contain a robust audit and logging system to ensure policy and regulatory compliance.

Collaboration functions, to include branching and version control. To allow for experimentation without requiring a separate development environment, the NAIRR platform should allow researchers to create, manage, and merge branches of their datasets, data pipelines, and models. This will not only accelerate and simplify the research and experimentation process, but allow for increased accountability and traceability of models, datasets, and data transformations within the NAIRR platform (e.g., ability to see when and where a new branch was created). Models and datasets could be directly compared in the platform to evaluate new techniques and identify sources of inaccuracy and bias. If a significant improvement in model accuracy, model fairness, or other metrics is validated within a branch, then researchers can share that back with the AI community, and leave annotations explaining the new branch. As researchers contribute to the shared data environment, the NAIRR will become the high-quality, validated, and collaborative data environment for AI model training and development.

User-friendly environment. To ensure the broadest possible participation within the research community, the user environment in the NAIRR should allow for easy manipulation of data even by non-technical users, and should include user-friendly security, compliance, and audit features for platform administrators. Given the existing technical barriers to entry within the AI research community, ensuring that the NAIRR is an inviting platform is critical to achieving the Task Force's goal of expanding the AI research community and democratizing access to AI tools.

Openness and interoperability. To achieve its full potential, the NAIRR must be compute-, storage-, and data format-agnostic. This ensures maximum flexibility for the government and broad participation from the research community. In compliance with data security and access controls, NAIRR users should be able to export their data, the code and logic in the data processing pipeline, the code responsible for building and running AI models, and the analysis code. This capability is critical to operationally deploying AI models outside the NAIRR or continuing research on external platforms using the same datasets and models.

1.E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets.

Based on our experience, successful data access and dissemination are linked with data infrastructure, including security and access controls, and stakeholder trust in the data infrastructure. We suggest that the NAIRR address barriers to sharing of high-quality datasets by avoiding a one-size-fits-all, all-or-nothing approach. Taking a more nuanced approach is likely to engender trust and confidence by government data holders and make them more amenable to sharing their data. While there likely are more than three relevant data categories, the examples below demonstrate the different levels of access control and data sharing policies the NAIRR should be able to accommodate:

- **Fully public data sources** with non-sensitive data such as ImageNet, Data.gov, and FBI Crime Data Explorer.

- **Enriched or annotated data sources** that may begin as public data sources or lightly-controlled data sources, but warrant policy and tooling to control their dissemination when enriched or annotated.
- **Combined data sources** formed by merging multiple controlled data sources or integrating researchers' external data sources with NAIRR-provided data. When combined, these data face many of the same issues as enriched or annotated data, and necessitate the option to evaluate them and control them appropriately.

The primary barriers to sharing more restricted, high-quality government datasets include:

Trust and Buy-In. Government agencies are less likely to share data if they cannot guarantee proper use and controlled dissemination. Technical infrastructure that enables secure data enclaves for multiple trusted partners to share the same data infrastructure can build trust and collaboration while ensuring security. See "Security" below and "Enable Multi-Stakeholder Engagement" on page 8.

Incompatible formats. The NAIRR platform should be interoperable to ensure that the widest possible selection of data can be ingested into the NAIRR platform for study. Additionally, the NAIRR should be able to automatically build and maintain live connections to relevant government source systems, so that models can be continuously tested on newly-generated data.

Security. To balance research security with inclusivity and flexibility, the NAIRR should include granular security and access control capabilities to allow for collaboration while ensuring compliance with privacy and civil liberty policies and regulations. A flexible framework for users to provide justifications to explain or qualify critical or risk-laden steps in their research workflows (e.g., data upload or export) should also be considered to reinforce propriety of user interactions with the NAIRR. Ensuring that the platform contains user-friendly and comprehensive administrator and management tools will also augment the ability for the research community to interact with government datasets.

Usability. Security control design can profoundly affect the usability and power of the platform for AI research. For example, Row-Based Access Controls, the ability for security markings and user permissions to be applied to each individual row of data, unlock many new opportunities for collaboration and platform accessibility. Instead of requesting access to an entire dataset or data pipeline, researchers may be able to use much of the information in a dataset while being restricted from certain elements (such as those that contain PII or other sensitive fields). The ability to automatically engage with specific subsets of data using granular security and access controls greatly expands the utility of the NAIRR and its accessibility to historically disadvantaged AI researchers. Additionally, the NAIRR should include a suite of deidentification tooling to further enable selective revelation of data, in compliance with necessity and proportionality considerations for data sharing, while mitigating risks of over-exposure and unintended privacy or data protection concerns.

1.F. An assessment of security requirements.

Security controls and requirements must be integrated into the foundational software fabric of the NAIRR, and should be automatically applied to not only data, but also to transformations, models, and combined datasets downstream of the original data. Please see our discussion of security in our responses to 1.E. and Question 3.

1.G. An assessment of privacy and civil rights and civil liberties requirements.

The NAIRR should adhere to relevant legal precedent as well as commonly accepted standards of privacy, civil rights, and civil liberties. This should include not only developing policy and governance standards for the platform, but also providing technical infrastructure inside the platform to help researchers maintain compliance with relevant standards. Cultural awareness considerations should factor into research efforts as well, both to help contextual the application of AI research to the complex social environments in which technology is actually ultimately used, but also as a means of anticipating ways that normative considerations may drive future legal and regulatory requirements. We provide further details on this topic in our response to Question 3.

2. Which NAIRR capabilities and services should be prioritized?

We recommend prioritizing a NAIRR data infrastructure that promotes and democratizes AI research. Please see our response to Question 1.D. above for tactical recommendations of priority capabilities to include in the NAIRR platform.

3. How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI?

Effective AI follows from the right data, properly managed. Through effective data integration and management, the NAIRR's policy and technical considerations should reinforce principles of responsible AI, which in turn will position the NAIRR to secure the trust of the AI research community and government data owners. To build trust, the NAIRR must be designed to provide visibility across data, processes, stakeholder organizations, and algorithms. In addition, NAIRR should provide mechanisms and toolkits to collect metadata on, understand, and remediate issues in underlying datasets that may contribute to forms of statistical or algorithmic bias. To reinforce and protect an ethical and responsible approach to AI R&D in NAIRR, Palantir suggests the following guiding principles:

Problem Prioritization, Review and Selection. There are certain problems that do not avail themselves of AI interventions. With the NAIRR providing a platform to democratize access to powerful tools and data, how the NAIRR evaluates and selects AI projects and datasets will be fundamental to the success of the program.

Methodically Assess and Address Sample Bias. The NAIRR can and should have the technical infrastructure to collect metrics about data quality so that it can assess and address sample and other forms of statistical bias. No matter how rigorous the design of an AI application is, if the data used for training and development is irredeemably flawed, the system output will almost certainly be compromised as well. This issue is even more pronounced where training data reflects systemic or other institutionalized forms of bias. As such, implementation of any AI system requires a clear assessment of the fidelity, quality, and representativeness of the data upon which its models are built and trained.

Ensure Auditability. To reinforce principles and practices of ethical AI, the NAIRR must provide accountability through automated traceability and auditing. In-platform system records and audit logs should be captured to allow for post hoc outcome analysis and documentation. The NAIRR is also in an ideal position to leverage the data collected in the platform to direct future research towards fundamental data or modeling gaps.

Ensure appropriate use with access controls. The platform should also include rigorous access controls that incorporate regulatory, legal, and normative expectations to provide the desired level of transparency, auditability, etc. Further details on our suggested approach to access controls are in our response to Question 1.D.

Enable Multi-Stakeholder Engagement. A critical operational principle for ethical AI usage is broad stakeholder involvement in building, deploying, and overseeing the NAIRR, as well as input in specific research selection and visibility into research outcomes. Multi-stakeholder engagement—particularly focused on communities impacted by AI—is critical to validating the NAIRR's mission and building trust across the community of interested and affected parties. However, incentivizing multi-stakeholder engagement also requires trust and the guarantee that partners can maintain a level of security and control over their data to ensure proper use. Relating to the themes of data quality and transparency articulated elsewhere in our response, we suggest that the NAIRR's technical infrastructure should be designed to enable multi-stakeholder engagement. This could include, but is not limited to:

- Infrastructure that supports secure enclaves within the same data environment.
- Collaboration and access control tools to allow the formation of data and research consortia.
- Visualizations attached to datasets to convey demographic composition intuitively.
- Built-in, automated compliance review workflows for sensitive data and research.
- Modeling results demonstrated with visualizations to promote understanding across various technical and non-technical stakeholders.

4. What building blocks already exist for the NAIRR?

While many potential building blocks exist around the government, such as ethical frameworks, the most important consideration for the NAIRR Task Force is to avoid duplicating efforts already taken by the public and private sectors. The government should leverage the billions of dollars of annual research and development currently undertaken by the private sector by engaging in public-private partnerships (see our response to Question 5 below for further details). For example, rather than building a new data infrastructure from scratch, NAIRR could partner with industry to use an off-the-shelf, continuously upgraded, open, cloud-agnostic software solution. With a technical infrastructure in place to allow for continuous delivery of new software features, the government would be able to not only maintain, but improve the NAIRR over time without major additional costs. We provide exemplars of private-public partnerships that can be used as models or building blocks in our response to Question 5 below.

5. What role should public-private partnerships play in the NAIRR?

Public-private partnerships are fundamental to ensuring the NAIRR becomes an effective shared research infrastructure, facilitating access for researchers and students to computational resources, high-quality data, educational tools, and user support. Recently, various U.S. Federal agencies leveraged innovative technologies to centralize and operationalize disparate COVID-19 data quickly, as well as a provide platforms to enable a unified response to the pandemic. Similarly, the NAIRR will rely on public-private partnerships to empower a broader swath of researchers and students to perform cutting edge AI research. Additionally, the advances gained through NAIRR—while funded by Government—will provide benefits across the public and private landscape.

Some relevant capabilities and outcomes across public-private health partnerships include:

- Industry provides intuitive data integration and analysis software to lower technical barriers and promote broad engagement by making datasets reusable, referenceable, and auditable.
- Government and Industry partner to consistently implement FAIR data principles, which could be similarly applied at the NAIRR, ensuring data is findable, accessible, interoperable, and reusable.
- Industry provides engineering data quality and security for Government research.

Palantir would also like to provide a few additional examples of our partnerships with the public sector to solve critical challenges through developing operational data assets.

U.S. National Institutes of Health (NIH) National COVID Cohort Collaborative (N3C) Consortium. Palantir software enables research across a variety of formerly disparate datasets and capabilities for the NIH N3C data enclave by automatically tracking data lineage of more than 5,000 data transformations across 65 sites, performing rigorous and automated data quality checks for trustworthy, research-ready data, harmonizing data at scale across more than 9.1 billion total records and over 3,000 users, and providing a collaborative workspace that has underpinned more than 30 scientific publications. Specific machine learning (ML) methods have been applied at N3C to improve understanding and response to COVID-19 and long COVID, including:

Analyzing large datasets of clinical and demographic data to understand correlation between demographic characteristics and increased clinical severity of COVID

Leveraging ML to predict clinical severity and risk factors over time across a study of nearly 2 million patients and 34 medical centers nationwide

U.S. Health and Human Services (HHS). Palantir software is the data infrastructure supporting the Department of Health and Human Service's Protect Platform. The Protect Platform assists the HHS and its partners to execute a holistic government response to fight the COVID-19 pandemic and protect the public's health, including providing a data infrastructure for analyzing public health data, such as the effect of school mask mandates on pediatric COVID transmission.

UK National Health Service (NHS). At the outset of the global COVID-19 pandemic in March 2020, the NHS deployed Palantir software to help determine how to best distribute life-saving equipment, including PPE and ICU consumable items. The NHS uses Palantir software to bring together over 150 datasets from across the NHS and partner organizations to enable a unified data foundation and single source of truth. Sources include hospital supply chain, epidemiological, staffing, atmospheric, emergency call center, and COVID-19 test result data.

Beyond Palantir's direct experience, the NAIRR should look to other Federal programs striving to promote equitable access to AI tools and datasets. This includes NIH's nascent Artificial Intelligence/Machine Learning Consortium to Advance Health Equity and Researcher Diversity (AIM AHEAD) program.

CONCLUSION & KEY RECOMMENDATIONS

The Task Force can proactively craft processes and policies that will sustain the NAIRR over the long-term. Prioritizing data quality and policies that incentivize and promote flexibility and a constantly improving technical infrastructure will empower NAIRR with a sustainable program to harness the power of the U.S. technology sector to benefit the AI research community. We have summarized our key recommendations in the chart below:

Theme	Recommendation
Infrastructure & Data Foundation	• Ensure researchers can easily discover training datasets, understand the data's provenance, and transparently and securely modify or annotate datasets in a way that benefit other researchers. • Provide tooling to allow researchers to contribute models, libraries, and updated datasets back to the NAIRR research community. • Implement privacy- and security-enhancing controls directly into the NAIRR data infrastructure, including audit logging capabilities. • Utilize open data, model formats, and documented APIs to allow researchers to export and import (while respecting permissions) to their preferred development environment. • Emphasize collaborative capabilities within the NAIRR to include data discovery, data provenance and transparency, as well as branching and versioning of data sets.
Responsible AI through Technical, Governance, and Cultural Awareness	• Adopt common standards around AI Privacy and Civil Liberties, Bias, and Ethics and enforce them through a combination of a) technical infrastructure to automate enforcement, b) incentivize and reward best practices, and c) cultural and discipline-specific frameworks to contextualize AI research and to guide determinations. • Provide researchers with an environment and tools to assess model performance both quantitatively and holistically. • Use NAIRR implementation to demonstrate effective U.S. global leadership for AI through ethical and responsible conduct. • Drive trust in NAIRR through data security and access controls, carrying out periodic security audits, providing visibility across data and data set transformations, as well as ensuring transparency of active user groups and research efforts. • Include diverse stakeholders (governmental, civil society organizations, researchers, industry, etc.) in periodic evaluations of NAIRR's efficacy and responsible employment.
Problem Prioritization through Resource Allocation	• Collect and analyze data about AI research focus areas; utilize incentives, including resource allocation, to shape those focus areas. • Utilize adjacent governmental organizations, such as the National Science Foundation and the White House Office of Science and Technology, to help identify under-researched areas.