

PUBLIC SUBMISSION

Received: May 29, 2025 Tracking No. nba-1svk-fwie Comments Due: May 28, 2025 Submission Type: Web
--

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0302
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Google

General Comment

Please see attached

Attachments

Google Comments re OSTP AI RandD RFI



**Response to the National Science Foundation's and
Office of Science & Technology Policy's Request for Information
on the Development of a 2025 National Artificial Intelligence (AI)
Research and Development (R&D) Strategic Plan**

90 Fed. Reg. 17835 (April 29, 2025)

Docket No. NSF-2025-OGC0001

May 29, 2025

AI is the defining technology of the 21st century, and it is central to America's economic prosperity, scientific leadership, and national security. While America currently leads the world in AI, this lead is not assured. Below, we set out three priorities for the National AI R&D Strategic Plan: (1) make strategic investments to secure U.S. AI leadership, (2) accelerate scientific discovery, and (3) strengthen security and resilience.

1. Make strategic investments to secure U.S. AI leadership

Invest in energy generation and modernize grid technology to support AI development

As we detailed in [our comments on the AI Action Plan](#) and our recent report, [Powering a New Era of American Innovation](#), ensuring a sufficient new energy supply to build out AI infrastructure is critical for maintaining U.S. AI R&D leadership. In addition to pursuing policy and permitting reforms to ensure adequate electricity supply to meet data center needs, the U.S. government should continue to fund R&D into novel power generation solutions; grid modernization, including optimization with AI; and worker education.

Make targeted research investments in key areas of national interest

It is incumbent upon the U.S. government to continue to prioritize support for foundational science in areas such as **algorithm & chip design, advanced manufacturing, and quantum technologies**, to ensure that the country can keep pace with competitors.

Invest in AI education and training

The new Strategic Plan should launch a comprehensive initiative to educate America's science students and scientific workforce on new AI technologies, supporting the uptake of [AI as the next scientific instrument](#). One pressing need is to fund and incentivize shorter, more tactical AI training programs and [fellowships](#) for scientists and research leaders.

Boost public sector innovation through talent exchanges and interoperability

The U.S. could champion **flexible talent mobility initiatives**. These could involve frameworks where AI researchers and engineers from the private sector undertake defined “tours of duty” or collaborative research projects within government agencies. Such programs could provide government agencies with direct exposure to cutting-edge techniques and talent. Additionally, to drive innovation in public services, the government should deploy standards-based technology where possible, at all layers of the stack, to **enable interoperability and data portability**.

2. Accelerate scientific discovery

Unlock the power of data and put advanced models in scientists’ hands

America needs to make **more compute capacity, high-quality data, AI models, and software tools available to more researchers**. This availability could be complemented by new public-private partnerships to get frontier AI labs’ models into the nation’s scientists’ hands as quickly as possible, including trusted tester programs that enable rapid feedback to improve the models’ utility for scientific research.

The U.S. government could systematically marshal its infrastructure and resources to help collect, complete, and make accessible key scientific data. For example, federal funding should **prioritize a series of ‘data stocktakes,’** where expert teams rapidly map the state of data in priority disciplines and application areas.

Balanced copyright rules, such as fair use, have been critical to enabling AI systems to learn from prior knowledge and publicly available data. They are essential for accelerating AI in science, particularly for applications that sift through scientific literature for insights or new hypotheses. The National AI R&D Strategic Plan should **champion and defend balanced copyright principles, especially fair use, as essential infrastructure for research**.

Science is a global enterprise, requiring international collaboration and access to global datasets. Disruptions to cross-border data flows and data localization mandates threaten to isolate U.S. researchers, reduce the overall data pool available for R&D, and diminish the quality and accuracy of AI models and tools. A key component of the U.S. AI R&D strategy must be to preserve and champion policies that enable **international data transfers**.

Identify critical challenges and design new incentives to drive progress

We recommend that the Strategic Plan identify a **list of priority scientific domains** and challenges to support economic development, health, and competitiveness goals. Federal funders could then incentivize progress, interdisciplinary collaboration, and public-private partnerships in a focused way through **innovative funding mechanisms**, such as public challenges, prizes, [advance market commitments](#), and fast grants.

Policymakers should also explore ways to accelerate the creation of new scientific evaluations. One idea would be to pilot a **dedicated office for Targeted Evaluations for Long-term Objectives in Science (TELOS)**. TELOS would commission compelling AI evaluations and competitions across the ecosystem. This demand-pull mechanism would provide clear, publicly legible metrics, serving as a powerful “carrot” for researchers.

Build evidence to adapt science for the age of AI

The new Strategic Plan could **invest in research on the use of AI in science** to highlight new opportunities for institutional innovation. In particular, scientists and policymakers have [explored only a small fraction](#) of the possible approaches to organizing and executing science research. AI provides a welcome forcing function to experiment with [new types of scientific institutions](#), and there will be a premium for countries that experiment the fastest.

3. Strengthen security and resilience

Strengthening our collective resilience against misuse and improving the reliability of AI systems are major technical and governance R&D challenges that must be addressed to advance security and adoption.

R&D for [Frontier Security Frameworks](#)¹

The U.S. government has specialized expertise and resources that can help accelerate and refine frontier security frameworks. We recommend additional efforts in the following areas:

Threat modeling: Foundational research into **threat modeling**—analysis of plausible pathways stemming from powerful capabilities of frontier models—is needed to prepare for novel misuse scenarios, such as AI-assisted cyberattacks or biological threats. Such modeling by frontier labs will be most effective when it draws on relevant deep domain expertise, often found in national security agencies.

Dangerous capability evaluations: Investment is needed to design new suites of dangerous capability evaluations. It’s particularly important for frontier AI labs and the U.S. government to collaboratively evaluate the capabilities of frontier models in areas where it has **unique expertise** and real-world threat intelligence, such as national security; chemical, biological, radiological, and nuclear issues; and cybersecurity threats.

Technical mitigations: Public R&D can help address two categories of mitigations: **security mitigations** intended to prevent the exfiltration of model weights, and **deployment mitigations** intended to counter the misuse of critical capabilities.

¹ See for example: Google’s updated [Frontier Safety Framework](#)—a set of protocols that aims to anticipate and address severe risks that may arise from powerful capabilities of foundation models; Anthropic’s [Responsible Scaling Policy](#); and OpenAI’s [Preparedness Framework](#).

Common frameworks and standards: In partnership with key industry actors, such as the Frontier Model Forum, public R&D could help **ensure that global standards are led by the U.S.** and build on U.S. industry best practice. These are vital for consistent assessment, preventing race-to-the-bottom dynamics and promoting adoption of U.S. technology.

Address fundamental safety & security research challenges

Public R&D can also drive progress on high-risk, high-reward fundamental research bets—including some that may pay out over longer periods of time—which could provide substantially stronger security and resilience as well as accelerate AI adoption. For example:

Agents & advanced security: R&D could explore agent identity and access controls; defenses against adversarial attacks on agents, such as prompt injections; least privileged principles which enable general purpose utility; and secure design for agent architectures.

Interpretability: In addition to promising recent advances in [interpretability techniques](#), more fundamental research is needed to gain a better mechanistic understanding of how frontier models work. Benefits range from enabling greater deployment of AI in high-stakes settings, to more trust in alignment techniques such that we can be sure an AI explanation of its answer faithfully represents its internal reasoning.

R&D for stronger societal resilience

Public R&D is also essential for equipping the U.S. government with advanced AI capabilities to address complex national security and societal challenges, while leveraging its unique mandates and expertise. These challenges include **bolstering cybersecurity** using AI capabilities and **strengthening biosecurity** against novel threats, each with opportunities for collaboration across the public and private sectors, as well as with allies.

We support public investment into aspirational **benchmarks that aim to drive innovation in areas of national strategic interest**. Many across the ecosystem play a role: government agencies outlining their urgent needs and providing domain expertise, AI developers innovating to address those needs, and the research community offering rigorous analysis to verify performance. We recommend the Administration champion an initiative that convenes key agencies (e.g., those in defense, intelligence, and diplomacy), national laboratories, academia, and industry to jointly develop these benchmarks and track progress.

As a longstanding leader in AI R&D, Google is committed to responsibly realizing the immense benefits of AI and supporting America's role as the world champion in AI innovation. We welcome the Administration's focus on this issue; with the right policy frameworks, America can look forward to an AI-powered golden era of opportunity.

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.