Google

Response to the National Science Foundation's and Office of Science & Technology Policy's Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

90 Fed. Reg. 9088 (Feb. 6, 2025)

Docket No. NSF_FRDOC_0001

March 13, 2025

Executive Summary

The potential of artificial intelligence is nearly unlimited, and we're already seeing how it can revolutionize healthcare, accelerate scientific discovery, and transform our economy for the better.[1] But a nation's ability to harness AI's enormous benefits requires the right policy frameworks.

Google welcomes the Trump Administration's goal of developing a plan to "sustain and enhance America's global AI dominance."[2] While America currently leads the world in AI—and is home to the most capable and widely adopted AI models and tools—our lead is not assured. As Vice President Vance urged, we must "catch lightning in a bottle" and unlock AI's potential.[3] To do that, we recommend focusing on three key areas to secure America's position as an AI powerhouse and support a golden era of opportunity:

1. **Invest in AI:** The Administration can take decisive actions to supercharge U.S. AI development, including:
    a. Coordinated federal, state, local, and industry action on policies like transmission and permitting reform to address surging energy needs, an essential part of expanding AI infrastructure.

---

[1] Yossi Matias, *Advancing healthcare and scientific discovery with AI*, Google (Mar. 4, 2025); Melissa Heikkilä, *Google DeepMind leaders share Nobel Prize in chemistry for protein prediction AI*, MIT Tech. Rev. (Oct. 9, 2024); McKinsey Digital, *The economic potential of generative AI: The next productivity frontier* (June 14, 2023).

[2] The White House, *Removing Barriers to American Leadership in Artificial Intelligence* (Jan. 23, 2025).

[3] The American Presidency Project, *Remarks by the Vice President at the Artificial Intelligence Action Summit in Paris, France* (Feb. 11, 2025).

    b. Balanced export controls that protect national security while enabling U.S. exports and global business operations.
    c. Continued funding for foundational AI research and development, streamlined access to computational resources for researchers, and public-private partnerships with national labs to advance research.
    d. Pro-innovation federal policy frameworks that preserve access to data for fair learning, advance a risk-based approach to AI applications based on existing regulations, and preempt a chaotic patchwork of state laws on frontier AI development.
2. **Accelerate and Modernize Government AI Adoption:** The federal government can lead by example through AI adoption and deployment, including implementing multi-vendor, interoperable AI solutions and streamlining procurement processes for new technologies.
3. **Promote Pro-Innovation Approaches Internationally:** The U.S. needs to pursue an active international economic policy to advocate for American values and support AI innovation internationally, including by:
    a. Championing market-driven and widely adopted technical standards and security protocols for frontier models, building on the Commerce Department's leading role with the International Organization for Standardization.
    b. Working with industry and aligned countries to develop tailored protocols and standards to identify and address potential national security risks of frontier AI systems.
    c. Combating restrictive foreign AI barriers that hinder American exports and innovation, while simultaneously promoting pro-innovation AI policies and establishing strong digital trade rules in future trade agreements.

We are in a global AI competition, and policy decisions will determine the outcome. A pro-innovation approach that protects national security and ensures that everyone benefits from AI is essential to realizing AI's transformative potential and ensuring that America's lead endures. Google is committed to doing its part, including by working with the federal government to ensure the success of the AI Action Plan. Our detailed recommendations follow.

1. Invest in AI

Like any multi-use technology, AI can be misused by bad actors, but it also promises to greatly improve our lives. For too long, AI policymaking has paid disproportionate attention to the risks, often ignoring the costs that misguided regulation can have on innovation, national competitiveness, and scientific leadership—a dynamic that is beginning to shift under the new Administration. Sustaining this momentum will require action in four areas:

A. Advance energy policies needed to power domestic data centers.

A potential lack of new energy supply is the core constraint to expanding AI infrastructure in the near term. Both training and inference computational needs for AI are growing rapidly. Compute requirements for training have historically doubled every six months, and inference compute needs are expected to increase by orders of magnitude in the coming years. While we are seeing significant efficiency improvements, widespread AI adoption may still result in large increases in electricity requirements, with projections of AI datacenter power demand rising by nearly 40 GW globally from 2024 to 2026.[4] Current U.S. energy infrastructure and permitting processes appear inadequate to meet these escalating needs.

The U.S. government should adopt policies that ensure the availability of energy for data centers and other growing business applications that are powering the growth of the American economy. This includes transmission and permitting reform to ensure adequate electricity for data centers coupled with federal and state tools for de-risking investments in advanced energy-generation and grid-enhancing technologies. Other key actions to meet new electricity load growth include improvements in electricity system planning, incentives for utilities to use existing infrastructure more efficiently, greater integration of regional electricity grids, and workforce development in building trades underpinning energy infrastructure.

B. Adopt balanced export control policies.

Export controls can play an important role in supporting national security, but only if they are carefully crafted to support legitimate market access for U.S. businesses while targeting the most pertinent risks. AI export rules imposed under the previous

[4] Dylan Patel et al., *AI Datacenter Energy Dilemma – Race for AI Datacenter Space*, Semianalysis (Mar. 13, 2024).

Administration (including the recent Interim Final Rule on AI Diffusion)[5] may undermine economic competitiveness goals the current Administration has set by imposing disproportionate burdens on U.S. cloud service providers. While we support the national security goals at stake, we are concerned that the impacts may be counterproductive and plan to submit a more detailed analysis of the AI Diffusion rule by the May 15 comment deadline.

The government will need to craft export controls carefully to avoid creating undue competitive disadvantages for U.S. companies. The U.S. government should adequately resource and modernize the Bureau of Industry and Security (BIS), including through BIS's own adoption of cutting-edge AI tools for supply chain monitoring and counter-smuggling efforts, alongside efforts to streamline export licensing processes and consideration of wider ecosystem issues beyond limits on hardware exports. Effective enforcement requires robust international engagement to maximize global compliance. And export controls are most impactful when coupled with a proactive strategy of domestic energy and infrastructure development to maintain a durable competitive advantage.

C. Accelerate AI R&D, streamline access to computational resources for researchers, and incentivize public-private partnerships with national labs.

Long-term, sustained investments in foundational domestic R&D and AI-driven scientific discovery have given the U.S. a crucial advantage in the race for global AI leadership. Policymakers should significantly bolster these efforts—with a focus on speeding funding allocations to early-market R&D and ensuring essential compute, high-quality datasets, and advanced AI models are widely available to scientists and institutions.[6] Lowering barriers to entry will ensure that the American research community remains keenly focused on innovation rather than struggling with resource acquisition. The government should also continue investments to identify and prioritize the most important unsolved challenges in the physical and life sciences (e.g., via federal prize challenges and competitions), focusing on how AI-driven approaches can help fuel scientific breakthroughs in areas of critical national interest.

---

[5] *See* *Framework for Artificial Intelligence Diffusion*, 90 Fed. Reg. 4544 (Jan. 15, 2025).
[6] Google, *A Policy Framework for Building the Future of Science with AI* (Feb. 2025).

Policymakers should move quickly to further incentivize partnerships with national labs to advance research in science, cybersecurity, and chemical, biological, radiological, and nuclear (CBRN) risks. The U.S. government should make it easier for national security agencies and their partners to use commercial, unclassified storage and compute capabilities, and should take steps to release government datasets, which can be helpful for commercial training.

D. Craft a pro-innovation federal framework for AI.

(i) Support federal legislation that prevents a patchwork of laws at the state level, especially for frontier AI development.

The Administration should ensure that the U.S. avoids a fragmented regulatory environment that would slow the development of AI, including by supporting federal preemption of state-level laws that affect frontier AI models. Such action is properly a federal prerogative and would ensure a unified national framework for frontier AI models focused on protecting national security while fostering an environment where American AI innovation can thrive. Similarly, the Administration should support a national approach to privacy, as state-level fragmentation is creating compliance uncertainties for companies and can slow innovation in AI and other sectors.

(ii) Ensure industry has access to openly available data that enable fair learning.

Three areas of law can impede appropriate access to data necessary for training leading models: copyright, privacy, and patents.

Copyright. Balanced copyright rules, such as fair use and text-and-data mining exceptions, have been critical to enabling AI systems to learn from prior knowledge and publicly available data, unlocking scientific and social advances. These exceptions allow for the use of copyrighted, publicly available material for AI training without significantly impacting rightsholders and avoid often highly unpredictable, imbalanced, and lengthy negotiations with data holders during model development or scientific experimentation. Balanced copyright laws that ensure access to publicly available scientific papers, for example, are essential for accelerating AI in science, particularly for applications that sift through scientific literature for insights or new hypotheses.

Privacy. Balanced privacy laws that recognize exemptions for publicly available information will avoid inadvertent conflicts with AI or copyright standards, or other impediments to the development of AI systems. A federal privacy regulatory framework should define categories of publicly available data and anonymous data that are treated differently than personally identifying data. Federal regulations can also encourage the use of AI-powered privacy-enhancing technologies to help protect Americans' data from malicious actors.

Patents. The Administration should improve and maintain access to the U.S. Patent and Trademark Office's Inter Partes Review program to permit efficient review of AI patents granted in error. The U.S. has seen tremendous growth in the patenting of AI in recent years.[7] Many of these patents are held by American companies like Google, but a growing percentage are held by entities based outside of the U.S., including in China.[8] In the last year, China's overall U.S. patent grants grew by over 30%, more than any other country.[9] With the increasing number of patent applications filed at the Patent and Trademark Office and the limited time available for reviewing those patent applications, mistakes are inevitable. According to one study, the agency's error rate may be nearly 40% for software-related technologies.[10] The rise of the first computers and then the internet saw a flood of patent applications for traditional functions simply performed "on a computer" or "via the internet." To avoid a similar phenomenon around functions performed "with AI," businesses need to be able to request agency assessments of a patent's validity through the Inter Partes Review process (when the high statutory bar is met). The agency should not reject meritorious requests based merely on agency-developed discretion (such as the *Fintiv* case), and needs to have continued staffing of its user-fee-funded Patent Trial and Appeal Board.[11] Otherwise, patents that were granted in error can be used by foreign entities to block and bottleneck American AI innovation, taking time and resources away from R&D, and subjecting highly sensitive technical information to discovery.

---

[7] Ayana Marshall, *AI Titans: Who's Dominating the Patent Universe*, Harrity (Mar. 11, 2024).
[8] Jack Caporal, *The Companies With the Most Generative AI Patents - and Why Investors Should Care*, Motley Fool (updated Mar. 9, 2025).
[9] IFI Claims, *2024 Trends and Insights* (last visited Mar. 12, 2025).
[10] Shawn P. Miller, *Where's the Innovation: An Analysis of the Quantity and Qualities of Anticipated and Obvious Patents*, 18 Va. J.L. & Tech. 1, 23 (2013).
[11] *See Apple Inc. v. Fintiv, Inc.*, IPR2020-00019 (Mar. 20, 2020).

(iii) Emphasize focused, sector-specific, and risk-based AI governance and standards.

Any regulation of AI applications should be proportional to relevant risks. Determining when, or if, to regulate requires context and a recognition of the unique challenges and opportunities in the specific domains where AI is used. Autocorrect features don't pose the same risks (or benefits) as healthcare applications deployed in an emergency room. To account for AI's context-dependent impacts, government regulation should be focused on specific applications, building upon existing sectoral rules and intervening directly only where demonstrably necessary.

Consensus technical standards and protocols can also play a critical role. As a baseline, regulations should align with recognized standards and support the development of standards and recommended practices; in many instances, establishing standards may be better than defining specific terms or thresholds in law or policy because they better keep pace with the technical state of the art. For example, standards and protocols can help ensure that privacy-enhancing technologies are implemented responsibly and in ways that make them accessible to businesses of all types and sizes, enable benchmarking, build trust, and protect Americans and their data.

(iv) Support workforce initiatives to develop AI skills and ensure American companies can hire and retain top AI talent.

AI is likely to contribute to important shifts in the future of work. While it can be easy to learn to use AI tools (since they can often teach the user how to use them), and the tools often benefit the least-skilled the most, the evolution of AI tools and deployment may still require a lifelong approach to education that gives all students and workers foundational AI skills.

This moment offers an opportunity to ensure that AI can be integrated as a core component of U.S. education and professional development systems. The Administration and agency stakeholders have an opportunity to ensure that access to technical skilling and career support programs (including investments in K-12 STEM education and retraining for workers) are broadly accessible to U.S. communities to ensure a resilient labor force.

In addition to workforce training and development, the ability of U.S. companies to access and retain top AI talent and expertise globally is essential and poses a known

challenge. Where practicable, U.S. agencies should use existing immigration authorities to facilitate recruiting and retention of experts in occupations requiring AI-related skills, such as AI development, robotics and automation, and quantum computing.

2. Accelerate and Modernize Government AI Adoption

To enable public sector organizations to fully benefit from the potential of cloud computing and AI, the government needs effective public procurement rules that foster innovation, ensure value for taxpayers, and promote a competitive and open market. The U.S. government, including the defense and intelligence communities, should pursue improved interoperability and data portability between cloud solutions;[12] streamline outdated accreditation, authorization, and procurement practices to enable quicker adoption of AI and cloud solutions; and accelerate digital transformation via greater adoption of machine-readable documents and data. We also encourage modernization of existing contracting processes to align with commercial procurement practices.

The federal government can also take advantage of opportunities to modernize procurement of emerging technology while reducing reliance on insecure legacy vendors. We propose lowering barriers to entry and growth through measures such as: (1) establishing reciprocity and harmonization for industry-approved certifications; (2) mandating re-use of existing authorizations and related materials to prevent duplication of effort; (3) facilitating investment in advanced threat detection; (4) instituting automated continuous monitoring methodologies; and (5) prioritizing open and market-based competition. Further, federal agencies should avoid implementing unique compliance or procurement requirements just because a system includes AI components. To the extent they are needed, any agency-specific guidelines should focus on unique risks or concerns related to the deployment of the AI for the procured purpose. U.S. decisionmakers might also consider policies to mandate interoperability throughout the entire technical stack and combat anticompetitive licensing and bundling practices. Doing so could also help ensure that government systems are not encumbered by known concentration risks of legacy technologies—many of which pose an unacceptable national security risk and cost more for the taxpayer.

---

[12] The Office of Management and Budget's (OMB's) 2024 AI Procurement Guidance outlined the importance of implementing multi-vendor, interoperable AI solutions. *See* Off. of Mgmt. & Budget, Exec. Off. of the President, OMB Memorandum M-24-18, *Advancing the Responsible Acquisition of Artificial Intelligence in Government* (2024).

Separately, policymakers should mandate open, non-proprietary data standards and APIs across all government cloud deployments, ensuring seamless interoperability and data portability to break down silos and enable AI-driven insights. As a part of this process, the current accreditation and procurement labyrinth should be replaced with a more agile, risk-based authorization process, drawing inspiration from commercial sector best practices to increase speed and accelerate the adoption of frontier AI and cloud solutions.

The Office of Science and Technology Policy (OSTP) and OMB can also issue guidance detailing more streamlined, automated, and responsive authorization processes for cloud services (including AI) under the Federal Risk and Authorization Management Program (FedRAMP); policies to advance greater reciprocity between agencies and their components; and a renewed approach to faster authorizations for AI services, which can have a transformative impact on federal agencies.

Policymakers should also consider measures to safeguard critical infrastructure and cybersecurity, including by partnering with the private sector. For example, pilots that build on the Defense Advanced Research Projects Agency's AI Cyber Challenge and joint R&D activities can help develop breakthroughs in areas such as data center security, chip security, confidential computing, and more. Expanded threat sharing with industry will similarly help identify and disrupt both security threats to AI and threat actor use of AI.

We recommend that the government continue its implementation of a multi-cloud and multi-model approach to national security use cases, which matches the most appropriate infrastructure and models to the agency, mission owner, and use case. We also recommend preserving existing risk-management guidelines covering AI use restrictions, minimum risk management practices for high-impact and federal personnel-impacting AI uses, and cataloging and monitoring AI use in the national security context.

3. Promote Pro-Innovation Approaches Internationally

To advance the widespread adoption of AI technologies both domestically and abroad, it is crucial to establish consistent, coherent, and interoperable frameworks and norms for AI development and deployment that reflect American values and interests.

Champion market-driven and widely adopted technical standards. Strong U.S. government support for standards based on American values will help keep foreign governments from imposing protectionist requirements that could stifle innovation, such as requiring duplicative pre-deployment testing to gain market access.

We encourage the Department of Commerce, and the National Institute of Standards and Technology (NIST) in particular, to continue its engagement on standards and critical frontier security work. Aligning policy with existing, globally recognized standards, such as ISO 42001, will help ensure consistency and predictability across industry.[13]

At the same time, rapid advances in frontier AI capabilities, including progress toward Artificial General Intelligence, highlight the need for the federal government to drive new efforts to ensure American leadership and national security. For the most capable frontier AI systems, the Administration should identify potential capabilities that could raise national security risks and work with industry to develop and promote standardized industry protocols, secure data-sharing, standards, and safeguards.

It is particularly valuable for the U.S. government to develop and maintain an ability to evaluate the capabilities of frontier models in areas where it has unique expertise, such as national security, CBRN issues, and cybersecurity threats. The Department of Commerce and NIST can lead on: (1) creating voluntary technical evaluations for major AI risks; (2) developing guidelines for responsible scaling and security protocols; (3) researching and developing safety benchmarks and mitigations (like tamper-proofing); and (4) assisting in building a private-sector AI evaluation ecosystem.

Building on the robust domestic approach outlined above, the U.S. government should work with aligned countries to develop the international standards needed for advanced model capabilities and to drive global alignment around risk thresholds and appropriate security protocols for frontier models. This includes promulgating an international norm of "home government" testing—wherein providers of AI with national security-critical capabilities are able to demonstrate collaboration with their home government on narrowly targeted, scientifically rigorous assessments that provide "test once, run everywhere" assurance. Reciprocity arrangements would enable other nations to acknowledge and accept home governments' evaluations,

[13] *See* ISO/IEC 42001 - Compliance | Google Cloud.

providing AI developers with appropriate market access without the need for additional government evaluations in those jurisdictions.

Articulate clear and differentiated obligations—where necessary—for the respective actors in the AI ecosystem. To the extent a government imposes specific legal obligations around high-risk AI systems, it should clearly delineate the roles and responsibilities of AI developers, deployers, and end users. The actor with the most control over a specific step in the AI lifecycle should bear responsibility (and any associated liability) for that step. In many instances, the original developer of an AI model has little to no visibility or control over how it is being used by a deployer and may not interact with end users. Even in cases where a developer provides a model directly to deployers, deployers will often be best placed to understand the risks of downstream uses, implement effective risk management, and conduct post-market monitoring and logging. Nor should developers bear responsibility for misuse by customers or end users. Rather, developers should provide information and documentation to the deployers, such as documentation of how the models were trained or mechanisms for human oversight, as needed to allow deployers to comply with regulatory requirements.

Avoid overbroad disclosure requirements. Policymakers should consider urging the use of model cards and technical reports—already an industry norm—in national and international fora to ensure that deployers and end users receive relevant information. The U.S. government should oppose mandated disclosures that require divulging trade secrets, allow competitors to duplicate products, or compromise national security by providing a roadmap to adversaries on how to circumvent protections or jailbreak models. Overly broad disclosure requirements (as contemplated in the EU and other jurisdictions) harm both security and innovation while providing little public benefit.

Notify users of AI-generated content in appropriate contexts. The U.S. government should support the further development and broad uptake of evolving multistakeholder standards and best practices around disclosure of synthetic media—such as the use of C2PA protocols, Google's industry-leading SynthID watermarking, and other watermarking/provenance technologies, including best practices around when to apply watermarks and when to notify users that they are interacting with AI-generated content. At the same time, the government should understand the limitations of such solutions—including the extent to which motivated

actors can strip out this information—and the need for cooperation among all players in the AI ecosystem to make progress on this issue.

Combat restrictive foreign AI barriers that hinder American businesses and innovation. Foreign regulatory regimes should foster the development of AI technology rather than stifle it. Governments should generally not impose regulatory checkpoints on the development of underlying AI models or AI innovation. Some governments are seeking to impose undue bureaucratic burdens on AI development and deployment, often in ways that would primarily affect U.S. companies. The U.S. government has a significant role to play in strengthening AI governance efforts and best practices by supporting innovation-friendly approaches and engaging foreign governments to deter efforts to impose measures that restrict AI development and deployment by U.S. and local companies. For example, OSTP and other federal stakeholders can consider bolstering and further resourcing interagency initiatives (including those undertaken by the State and Commerce Departments) that target and strengthen commercial diplomacy and promote exports of U.S. digital goods and services, including American AI. And the U.S. should advocate at the Organisation for Economic Co-operation and Development (OECD) and other fora for international AI frameworks that reflect U.S. values and approaches.

******

As a longstanding leader in AI research and development, Google is committed to responsibly realizing the immense benefits of AI and supporting America's role as the world champion in AI innovation. Our mission is to organize the world's information and make it universally accessible and useful, and our work on AI lies at the heart of that mission. We welcome the Administration's focus on this issue, and we agree that with the right policy frameworks, America can look forward to an AI-powered golden era of opportunity.[14]

[14] 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 AI Action Plan and associated documents without attribution.