

Networking and Information Technology Research and Development, National Coordination Office
National Science Foundation
Document No. 2025-02305
Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

March 15, 2025

Introduction

We recommend that the Office of Science and Technology Policy (OSTP) and the Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO) prioritize the following four areas of policy action in their AI Action Plan:

1. Promote open innovation as a strategic advantage for U.S. competitiveness.
2. Maintain U.S. AI leadership by promoting scientific innovation.
3. Craft evidence-based AI policy that protects Americans without stifling innovation.
4. Empower government leaders with resources and technical expertise to ensure a “whole-of-government” approach to AI governance.

1. Promote open innovation as a strategic advantage for U.S. competitiveness

Open innovation has been a cornerstone of the U.S. AI ecosystem, fueling collaboration, competition, and rapid experimentation. To maintain this advantage, the United States must protect, preserve, and promote access to open models rather than impose broad restrictions that could slow innovation and isolate domestic developers.

Promote open innovation: The United States has long led in AI innovation by fostering an ecosystem that encourages broad access to foundational research and technology.¹ This approach allows a diverse set of researchers—from startups to technology firms and academic institutions—to build on shared advancements, driving faster breakthroughs and greater technological resilience. Open foundation models, in particular, exemplify how open innovation accelerates progress. By making model weights widely available, these models enable greater customizability, foster transparency, and enhance competition. They allow researchers to fine-tune models for specific applications, improve scientific reproducibility, and scrutinize systems for potential risks, reinforcing both innovation and security. Ensuring continued access to open models strengthens the broader AI ecosystem and sustains American leadership in cutting-edge development.²

¹ Peter L. Singer, “Federally Supported Innovations,” *Information Technology & Innovation Foundation*, February 2014, <https://www2.itif.org/2014-federally-supported-innovations.pdf>.

² Sayash Kapoor et al., “Position: On the Societal Impact of Open Foundation Models,” *Proceedings of the 41st International Conference on Machine Learning*, July 2024, <https://dl.acm.org/doi/10.5555/3692070.3692998>; National Telecommunications and Information Administration, “Dual-Use Foundation Models with Widely Available Model Weights Report,” July 30, 2024, <https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report>.

To sustain this leadership, the federal government should support policies that preserve access to open foundation models while advancing safeguards against misuse. This includes funding research on secure model-sharing practices, developing technical benchmarks for responsible openness, and supporting technical solutions like improved tracking and security measures over rigid liability frameworks that are impractical for open model developers.

Avoid restrictive domestic policies: Restricting open development of AI would weaken U.S. competitiveness without meaningfully curbing the advancements of foreign adversaries. China’s AI ecosystem, for example, continues to evolve, with models like DeepSeek demonstrating how the open-source approach accelerates technological progress. China’s AI trajectory is not dependent on U.S. openness—its government and companies are already developing powerful models which are being adopted around the world. Broad restrictions would isolate U.S. developers and reduce the transparency of AI systems, limiting the ability to develop robust AI security measures and weakening the overall resilience of the innovation ecosystem. It would also result in more global dependencies on Chinese rather than U.S. open models. Instead of imposing sweeping restrictions, the federal government should pursue a balanced approach that supports open development while implementing targeted safeguards to mitigate national security risks and prevent misuse.

2. Maintain U.S. AI leadership by promoting scientific innovation

The United States has long maintained global leadership in AI through sustained public investment in foundational research. ImageNet, a federally funded project, for example, was a cornerstone of the deep learning revolution that powers modern AI.³ To continue this trajectory, the United States must commit to long-term investment in AI research, ensuring that innovation remains competitive, national security interests are protected, and scientific progress is not dictated solely by short-term commercial priorities.

Sustain federal funding for basic research: Federal funding has driven critical breakthroughs in deep learning and natural language processing—advancements that have shaped the modern AI landscape. Continued federal investment is essential to sustaining U.S. innovation and ensuring that AI systems remain both cutting-edge and aligned with democratic values. Public investment enables research that makes AI more efficient, accurate, and adaptable—advancements that drive scientific discovery and promote the responsible use of AI across industries. From developing next-generation AI architectures to enhancing robustness and interpretability, federally funded research plays a pivotal role in pushing the boundaries of innovation.

As China accelerates its state-backed AI research initiatives, the United States must maintain a strategic edge through sustained investment in its own AI ecosystem.⁴ Without robust federal funding, American

³ Jia Deng et al., “ImageNet: A Large-Scale Hierarchical Image Database,” *2009 IEEE Conference on Computer Vision and Pattern Recognition*, June 2009, <https://ieeexplore.ieee.org/document/5206848>.

⁴ “China Steps Up Support for Tech Sector as AI Enthusiasm Soars,” *Bloomberg*, March 2025, <https://www.bloomberg.com/news/articles/2025-03-06/china-steps-up-support-for-tech-sector-as-ai-enthusiasm-soars>.

researchers risk falling behind, weakening scientific discovery, shrinking the AI talent pipeline, and increasing reliance on private sector funding, which often prioritizes short-term applications over foundational innovation. History shows that academia, unbound by immediate profit incentives, has driven transformative innovations that industry would not have pursued alone—such as CRISPR reshaping genetics and early internet protocols laying the foundation for the digital economy.⁵ Federal research support ensures that AI development serves the broad national interests of the public.

Establish and fund the National AI Research Resource (NAIRR): Our AI innovation ecosystem—especially academia, startups, and public sector players—is facing challenges due to the high cost of compute resources and lack of access to government data. This lack of accessible infrastructure and data hampers research, slows innovation, and threatens U.S. leadership, especially as competitors such as China make substantial investments in state-backed AI resources.⁶ The NAIRR will provide critical compute and datasets to researchers beyond major technology firms, ensuring that academic institutions, startups, and governments can meaningfully contribute to AI development and deployment.⁷

For example, during the COVID-19 pandemic, the U.S. Department of Energy’s supercomputing resources were harnessed to accelerate drug discovery and understand the virus’s behavior, showcasing how focused compute allocation can lead to rapid, lifesaving advancements.⁸ Similarly, broader sharing of government data has exposed inefficiencies in key public systems, such as tax administration and mass adjudication, allowing researchers to develop evidence-based policy improvements.⁹ Fully funding and implementing NAIRR will strengthen U.S. leadership, accelerate scientific progress, and prevent AI development from becoming overly concentrated within a few dominant corporations.

⁵ See Leiner et al., “A Brief History of the Internet,” *ACM SIGCOMM Computer Communication Review* 39(5), 2009, <https://dl.acm.org/doi/10.1145/1629607.1629613>; Jennifer A. Doudna and Emmanuelle Charpentier, “The New Frontier of Genome Engineering with CRISPR-Cas9,” *Science* 346 (6213), 2014, <https://www.science.org/doi/10.1126/science.1258096>.

⁶ Ben Jiang, “China’s Capital Beijing to Provide State-Sponsored Computing Resources to AI Firms Amid ChatGPT frenzy,” *South China Morning Post*, May 16, 2023, <https://www.scmp.com/tech/policy/article/3220736/chinas-capital-beijing-provide-state-sponsored-computing-resources-ai-firms-amid-chatgpt-frenzy>.

⁷ Daniel E. Ho et al., “Building a National AI Research Resource,” *Stanford Institute for Human-Centered Artificial Intelligence and Stanford Law School*, October 2021, https://hai-production.s3.amazonaws.com/files/2022-01/HAI_NRCR_v17.pdf; National Artificial Intelligence Research Resource Task Force, “Strengthening and Democratizing the US Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource,” January 2023, <https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf>.

⁸ “Q&A with Under Secretary Paul Dabbar on the COVID-19 High Performance Computing Consortium,” U.S. Department of Energy, <https://www.energy.gov/articles/qa-under-secretary-paul-dabbar-covid-19-high-performance-computing-consortium>.

⁹ See research from the Stanford Regulation, Evaluation, and Governance Lab, <https://reglab.stanford.edu>, e.g., Peter Henderson et al., “Integrating Reward Maximization and Population Estimation: Sequential Decision-Making for Internal Revenue Service Audit Selection,” *Thirty-Seventh AAAI Conference on Artificial Intelligence*, 2023, <https://ojs.aaai.org/index.php/AAAI/article/view/25637>; Daniel E. Ho et al., “Quality Review of Mass Adjudication: A Randomized Natural Experiment at the Board of Veterans Appeals, 2003–16,” *The Journal of Law, Economics, and Organization* 35 (2), March 2019, https://dho.stanford.edu/wp-content/uploads/Ho_HandanNader_Ames_Marcus.pdf.

3. Craft evidence-based AI policy that protects Americans without stifling innovation

When crafted pragmatically, AI policy that is science- and evidence-based can mitigate serious harms without hampering private sector AI innovation.¹⁰ OSTP should pursue low-cost, high-impact AI safeguards that, by mitigating unintended consequences and holding companies to reasonable standards, will build trust among users and foster healthy competition among developers. Such policies will ultimately help increase the commercialization and adoption of AI, while strengthening U.S. leadership in AI.

Our research has shown there are several concrete policy levers that provide clear benefits while requiring relatively few resources to operationalize. These levers primarily address a persistent lack of reliable information about the capabilities and risks of different AI systems—a lack that remains a central impediment to effective AI governance.

Transparency requirements: Mechanisms to promote, mandate, and specify transparency in AI systems should be a priority. Transparency, for example, regarding AI developers’ business practices, the resources they use to build models, and the risks and limitations of their models, is crucial for both policymakers and the general public to better understand and trust powerful AI systems used by hundreds of millions of people.¹¹ It also enables policymakers to better design and enforce AI regulation. As improvements on our Foundation Model Transparency Index show, transparency can be relatively low cost to developers, and applying transparency standards broadly can incentivize a race to the top.¹²

Adverse event reporting: Policy that mandates or incentivizes the transparent reporting of harmful AI behavior—both concrete incidents (e.g., misdiagnosis by medical AI systems) and more abstract concerns (e.g., the generation of biological pathogens)—is also crucial.¹³ It would enable policymakers to aggregate information about adverse events and incidents arising from specific AI systems.¹⁴ Comparable reporting systems at the Food and Drug Administration and the Cybersecurity and Infrastructure Security Agency have shown that such a policy will greatly aid regulators as they monitor emergent risks and identify

¹⁰ Fei-Fei Li, “Now More Than Ever, AI Needs a Governance Framework,” *Financial Times*, February 8, 2025, <https://www.ft.com/content/3861a30a-50fc-41c9-9780-b16626a0d2e8>; Rishi Bommasani et al., “A Path for Science- and Evidence-based AI Policy,” <https://understanding-ai-safety.org/>.

¹¹ Rishi Bommasani et al., “The Foundation Model Transparency Index (October 2023),” *Center for Research on Foundation Models*, October 2023, <https://crfm.stanford.edu/fmti/October-2023/index.html>.

¹² Rishi Bommasani et al., “The Foundation Model Transparency Index (May 2024),” *Center for Research on Foundation Models*, May 2024, <https://crfm.stanford.edu/fmti/>.

¹³ The National Artificial Intelligence Advisory Committee (NAIAC), “RECOMMENDATION: Improve Monitoring of Emerging Risks From AI Through Adverse Event Reporting,” November 2023, https://ai.gov/wp-content/uploads/2023/12/Recommendation_Improve-Monitoring-of-Emerging-Risks-from-AI-through-Adverse-Event-Reporting.pdf.

¹⁴ Neel Guha et al., “AI Regulation Has Its Own Alignment Problem: The Technical and Institutional Feasibility of Disclosure, Registration, Licensing, and Auditing,” *George Washington Law Review* 92, December 2024, <https://www.gwlr.org/wp-content/uploads/2024/12/92-Geo.-Wash.-L.-Rev.-1473.pdf>.

trends that demand regulation or guidance, while requiring relatively few technical and institutional resources.¹⁵

Safe harbors for third-party research: The AI Action Plan should also prioritize policy that safeguards independent researcher access to AI models. Third-party evaluation is a cornerstone of efforts to reduce the substantial risks posed by AI systems, yet few companies actively protect or promote such research.¹⁶ Mandating that AI companies adopt legal and technical safe harbors for good-faith AI safety and trustworthiness research, while also securing a commitment from the Department of Justice not to criminally charge such activities,¹⁷ would protect important good-faith research and stress-testing. It would also significantly improve policymakers' and the general public's understanding of the risks of AI and hold developers accountable.

4. Empower government leaders with the resources and technical expertise to ensure a “whole-of-government” approach to AI governance

OSTP's AI Action Plan will only be successful if all federal agencies and departments have the necessary high-level leadership, resources, and technical expertise to implement it. Our work tracking and analyzing the state of implementation of a variety of AI-related executive actions released since 2019, beginning with the first Trump administration, has shown that to effectively govern AI, federal agencies must fill gaps in senior leadership, agency capacity, and technical talent, while also pursuing close collaboration with state policymaking and implementation.¹⁸

Ensure coordinated, high-level leadership: Federal agencies and the White House must drive government-wide AI governance efforts through coordinated, high-level leadership. In the past, leadership vacuums contributed to inconsistent implementation of executive actions on AI.¹⁹ OSTP should preserve the chief AI officer (CAIO) role and expand its mandate beyond overseeing compliance with risk management measures: CAIOs should be empowered to more holistically balance agency measures to advance AI innovation with safeguards that protect citizens from AI-related harms.²⁰ However, leaders will require adequate resources and personnel to achieve this.

¹⁵ Neel Guha et al., “The AI Regulatory Alignment Problem,” *Stanford Institute for Human-Centered Artificial Intelligence*, <https://hai.stanford.edu/policy/policy-brief-ai-regulatory-alignment-problem>.

¹⁶ Kevin Klyman et al., “Safeguarding Third-Party AI Research,” *Stanford Institute for Human-Centered Artificial Intelligence*, February 13, 2025, <https://hai.stanford.edu/policy/safeguarding-third-party-ai-research>.

¹⁷ See “Department of Justice Announces New Policy for Charging Cases under the Computer Fraud and Abuse Act,” U.S. Department of Justice, May 19, 2022, <https://www.justice.gov/archives/opa/pr/departments-justice-announces-new-policy-charging-cases-under-computer-fraud-and-abuse-act>.

¹⁸ “Tracking U.S. Executive Action on AI,” *Stanford Institute for Human-Centered Artificial Intelligence*, <https://hai.stanford.edu/policy/tracking-us-executive-action-ai>.

¹⁹ Christie Lawrence et al., “Implementation Challenges to Three Pillars of America's AI Strategy,” *Stanford Institute for Human-Centered Artificial Intelligence*, December 20, 2022, <https://hai.stanford.edu/policy/white-paper-implementation-challenges-three-pillars-americas-ai-strategy>.

²⁰ Jennifer Wang et al., “Assessing the Implementation of Federal AI Leadership and Compliance Mandates,” *Stanford Institute for Human-Centered Artificial Intelligence*, January 17, 2025, <https://hai.stanford.edu/policy/assessing-the-implementation-of-federal-ai-leadership-and-compliance-mandates>.

Provide adequate implementation resources: Agency leaders require the appropriate financial resources, digital infrastructure, and technical expertise to be able to develop and meaningfully execute strategic AI plans. For example, the Department of Commerce and the National Institute of Standards and Technology (NIST) require sustained investment in funding, personnel, and computing capacity to effectively conduct the measurement, evaluation, and standardization of AI models. Ensuring agencies have the most up-to-date software and computing resources is also crucial for attracting and retaining technical talent, which is critical to ensuring the successful implementation of AI policies.²¹

Strengthen internal technical expertise: Securing a robust pipeline of technical talent across all federal agencies must be a priority. This includes upskilling civil servants in their understanding of AI and using academic-agency partnerships to temporarily bring talent into the government.²² It is crucial that agencies develop internal talent to avoid over-reliance on external vendors to perform core AI governance oversight functions.

Further collaboration with state policymakers: With over 700 AI-related bills introduced in 2024 and more expected in 2025, states are taking the lead on AI governance.²³ While state-driven innovation is valuable, inconsistent policies can create uncertainty and operational challenges. The executive branch should facilitate two-way engagement with states—ensuring that states can benefit from sharing technical expertise, best practices, and economic and security considerations while also informing federal policymakers of on-the-ground challenges and innovations.

Federal agencies already collaborate with states on cybersecurity;²⁴ AI governance should follow suit. NIST can provide voluntary guidance and standardization frameworks, helping states craft policies that foster innovation while respecting state autonomy. Strengthening federal-state collaboration will ensure effective, evidence-based AI governance that balances innovation, security, and regulatory clarity.

We thank the OSTP, NSF, and NITRD NCO for the opportunity to share our views, which are based on our scientific research in these areas. Please email cm21@stanford.edu and dzhang105@stanford.edu with any comments or questions.

²¹ Daniel E. Ho, “Opportunities and Risks of Artificial Intelligence in the Public Sector,” *Testimony Presented to the US Senate Committee on Homeland Security and Governmental Affairs*, May 16, 2023, <https://hai-production.s3.amazonaws.com/files/2023-05/Daniel-Ho-Senate-Testimony.PDF>.

²² Mariano-Florentino Cuéllar et al., “Response to OMB’s Request for Comment on Draft Policy Guidance on Agency Use of AI,” November 30, 2023, <https://hai-production.s3.amazonaws.com/files/2024-01/Response-Stanford-RegLab.pdf>.

²³ Business Software Alliance, “2025 State AI Wave Building After 700 Bills in 2024,” October 22, 2024, <https://www.bsa.org/news-events/news/2025-state-ai-wave-building-after-700-bills-in-2024>.

²⁴ For example, the Department of Homeland Security (DHS) supports the Multi-State Information Sharing and Analysis Center (MS-ISAC), which serves as a central hub for state and local threat intelligence. MS-ISAC provides free cybersecurity resources, including network monitoring tools, to help states strengthen their digital infrastructure. See “Multi-State Information Sharing and Analysis Center,” Cybersecurity & Infrastructure Security Agency, <https://www.cisa.gov/resources-tools/services/multi-state-information-sharing-and-analysis-center>.



Sincerely,

Caroline Meinhardt

Policy Research Manager, Stanford Institute for Human-Centered Artificial Intelligence

Daniel Zhang

Senior Manager for Policy Initiatives, Stanford Institute for Human-Centered Artificial Intelligence

Rishi Bommasani

Society Lead, Stanford Center for Research on Foundation Models

Ph.D. Candidate, Stanford University

Jennifer King

Privacy and Data Policy Fellow, Stanford Institute for Human-Centered Artificial Intelligence

Russell Wald

Executive Director, Stanford Institute for Human-Centered Artificial Intelligence

Percy Liang

Senior Fellow, Stanford Institute for Human-Centered Artificial Intelligence

Director, Stanford Center for Research on Foundation Models

Associate Professor of Computer Science and (by courtesy) Statistics, Stanford University

Daniel E. Ho

Senior Fellow, Stanford Institute for Human-Centered Artificial Intelligence

Director, Stanford Regulation, Evaluation, and Governance Lab

William Benjamin Scott and Luna M. Scott Professor of Law, Professor of Political Science, and (by courtesy) Computer Science, Stanford University