NEW AMERICA

Technology & Democracy Programs

March 15, 2025

DOCKET NSF_FRDOC_0001
Re. National Science Foundation's Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

(Submitted via email and regulations.gov)

New America's Technology & Democracy Programs work to foster a sustainable digital future that advances equitable opportunity, innovation, fundamental rights, and participatory governance—where democracy, human rights, and the planet flourish.[1] We appreciate the opportunity to respond to the request for information regarding the development of an AI Action Plan, as required by Executive Order 14179 ("E.O. 14179").[2]

An AI Action Plan Should Build on Existing Foundations

E.O. 14179's foundational assumption is that the prior administration's AI policies hampered the U.S. government's and American companies' ability to compete globally. But the first Trump administration established key foundations on which the Biden administration built. An AI Action Plan should preserve these areas of bipartisan continuity over the last eight years if the United States is to retain its global competitive advantage. Much of the now-revoked E.O. 14110 was designed to promote the "human flourishing" that E.O. 14179 seeks to promote.

AI cannot enable human flourishing without the trust required for AI adoption. Trust in turn depends upon continued progress in ensuring AI systems are safe, effective, and fair. Our comments emphasize the ways in which maintaining our commitment to fairness, openness, and democratic values are at the core of U.S. global competitiveness in AI innovation.

---

[1] Technology and Democracy Programs, *New America*, https://www.newamerica.org/technology-and-democracy/about/.
[2] Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence, 90 Fed. Reg. 8741 (Jan. 31, 2025), https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence.

Trust in American public- and private-sector uses of AI can only be built on U.S. governance frameworks that balance innovation with democratic values and human agency. This trust is essential to global consumer confidence and adoption, as American companies face stiff market competition. Rigorous evaluation of AI systems for accuracy, fairness, bias, and safety are preconditions for producing AI systems that can outperform competitors.

A recent paper on improving evaluation science for generative AI powerfully defines the relationship between evaluation and trust: "The bridges we stand on, the medicine we take, and the food we eat are all the result of rigorous assessment. In fact, it is because of the rigor of the corresponding evaluation ecosystems that we can trust that the products and critical infrastructure surrounding us are performant and safe. Generative AI products are no exception to this reality."[3] A forthcoming AI Action Plan should thus emphasize for companies and consumers a continued commitment to rigorously evaluating public- and private-sector AI systems both before and after they are deployed.

Calls for a continued focus in AI governance on safety, fairness, and accountability are not about advancing politicized agendas. They *are* about careful vetting, protecting the American people, and ensuring global and consumer trust in U.S.-developed AI, all of which enjoy bipartisan support.[4] Indeed, the first Trump administration issued Executive Order 13960 focused squarely on these objectives and established important transparency practices that promoted them.[5]

Openness in the AI ecosystem is an important enabler of trust and accountability. It is also a building block of promoting innovation, competition, and safety. Encouraging openness in AI systems serves multiple aims: innovation, developing applications in the public interest, and transparency about both code and model governance that helps make AI systems safer, more effective, and less vulnerable to cyberattacks.

---

[3] Laura Weidinger and Hannach Wallach et al., "Toward an Evaluation Science for Generative AI Systems," Mar. 2025, https://arxiv.org/html/2503.05336v1.

[4] *See*, *e.g.*, Bipartisan House Task Force Report on Artificial Intelligence (Dec. 2024), https://republicans-science.house.gov/_cache/files/a/a/aa2ee12f-8f0c-46a3-8ff8-8e4215d6a72b/A163BDBF496ADA741F831E5BEBBCA06699B6AFF8CC34F4FDC4065BDA298295DF.ai-task-force-report-final.pdf (noting the imperative to "protect Americans from accidental and malicious uses of AI" and explaining that "[t]rust is a necessary component for the widespread adoption of AI by the public and private sectors in the United States.")

[5] Executive Order 13960, Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government, Dec. 8, 2020, https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government ("The ongoing adoption and acceptance of AI will depend significantly on public trust. Agencies must therefore design, develop, acquire, and use AI in a manner that fosters public trust and confidence while protecting privacy, civil rights, civil liberties, and American values, consistent with applicable law").

An AI Action Plan Should Respond to the Dangers of Technopoly and Technopolarization

It is clear that the ascendance of artificial intelligence (AI) presents both unprecedented opportunities and serious challenges for American technological leadership. As we navigate this rapidly evolving landscape, the United States must pursue governance strategies that balance multiple competing priorities: maintaining technological primacy, engaging global partners, addressing legitimate concerns about AI risks, respecting rights and liberties, and adapting to domestic political realities that favor fiscal restraint and skepticism of new international commitments.

We stand at a critical juncture where “technopoly”—the surrender of culture and governance to technological imperatives—threatens to become the default operating system for advanced economies. When powerful AI systems are developed primarily by a handful of companies and nations, the risk of technopolarization grows acute: a world divided between those who control AI systems and those who are merely subject to them. This concentration of power would not only entrench existing global inequities; it would also undermine long-term U.S. interests by creating fertile ground for competing powers to position themselves as technological liberators.

Policy Recommendations

We offer eight recommendations that leverage America's core strengths—its considerable state capacity in AI governance and scientific research, a dynamic private sector, world-class research institutions, and vibrant civil society—while acknowledging the necessity of global engagement. This approach emphasizes market-driven partnerships, multi-stakeholder networks, and strategic convening power that distributes rather than concentrates technological agency.

By incorporating multiple stakeholders and investing in partnerships in the Global South, we can build governance structures that resist both technopoly and technopolarization, as well as foster global consumer trust in American AI companies. This approach recognizes that the Global South represents not just markets for American innovation but essential partners whose perspectives are vital to developing AI systems that serve truly global needs rather than reinforcing existing power imbalances.

Recommendation 1: Develop AI governance standards informed by democratic values.
The United States should participate in developing governance frameworks that balance innovation with democratic values and human agency. Multi-stakeholder working groups can establish standards that preserve technological dynamism while ensuring

appropriate safeguards against concentrated power. The United States should prioritize developing metrics that evaluate AI systems on their contribution to user empowerment and expanded choice, not just technical performance. This approach directly addresses technopoly concerns and engenders global consumer trust by ensuring that technological governance is anchored in broader societal values.

Recommendation 2: Ensure continued U.S. leadership in AI standard-setting and evaluation by adequately resourcing NIST and the U.S. AI Safety Institute.
More specifically, deploying trustworthy AI requires effective assessments and the development of strong best practices and standards. The National Institute of Standards and Technology (NIST) is central to these efforts. For more than a century, NIST has provided trusted, impartial leadership in collaboratively developing well-defined technical standards and promoting innovation across a range of American industries.

As the complexity and scale of AI systems increase, NIST is uniquely qualified to help develop the technical standards, testing methods, and objective evaluation techniques required for effective AI governance frameworks. Some of NIST's most notable contributions to AI governance include developing the AI Risk Management Framework, establishing the U.S. AI Safety Institute (USAISI), and driving multi-stakeholder work on measurement science in AI safety. These painstaking efforts have built state capacity that serves American interests.

However, workforce reductions at NIST threaten to erode that American capacity and competitive advantage. News reports also indicate that NIST is no longer seeking scientific partners with expertise in "AI fairness," "safety," or "responsible AI" in favor of scientists committed to "removing ideological bias."[6] These actions also run contrary to the U.S. government's goal of maintaining U.S. leadership. Prioritizing AI safety, fairness, and accountability is not promoting a politicized agenda; it is about doing what is required to develop effective AI models and products that U.S. and international consumers are willing to use. Fairness, safety, and responsible AI are not just central to preserving E.O. 14179's goal of "human flourishing," but also advance economic competitiveness and national security. An AI Action plan should recognize NIST's (and USAISI's) importance to developing trustworthy AI.

[6] Will Knight, "Under Trump, AI Scientists are Told to Remove 'Ideological Bias' From Powerful AI Models," *Wired* (Mar. 14, 2025), https://www.wired.com/story/ai-safety-institute-new-directive-america-first/.

Recommendation 3: Preserve and expand governmental efforts to boost American capacity on AI, including via the National Artificial Intelligence Research Resource (NAIRR) Pilot.

The NAIRR Pilot has helped to marshal the collective innovative potential of the federal government, large technology companies, and U.S. scientific institutions. It continues a decades-long tradition of public-private-researcher cooperation that has fueled U.S. global leadership. NAIRR provides infrastructure that other stakeholders can use to build, understand, and use frontier models. It democratizes access to resources in ways that spur AI innovation and allow smaller players to compete with large tech firms. It is a critical part of maintaining competitiveness within the U.S. economy and enabling our private and nonprofits sectors' ability to produce unexpected, transformative innovation that leads the world. The administration should partner with Congress to permanently establish NAIRR, moving it beyond its pilot status.

Recommendation 4: Maintain the transparency and safety requirements of OMB Memorandum M-24-10.

The Office of Management and Budget should maintain the transparency and safety protections of Memorandum M-24-10 as it implements President Trump's E.O. 14179.[7] Two aspects of the memorandum are central to continued American leadership in AI innovation, effectiveness, trustworthiness, and safety: the AI use case inventories (along with related transparency provisions) and the minimum practices for rights-and safety-respecting AI.

Thirteen research and public interest organizations recently wrote to the OMB to emphasize that the repeal of E.O. 14110 does not require agencies to stop publishing use case inventories.[8] In fact, these components of M-24-10 have their roots in actions taken by the first Trump administration. During his first term, President Trump instructed agencies to "design, develop, acquire, and use AI in a manner that fosters public trust and confidence while protecting privacy, civil rights, civil liberties, and American values."[9]

---

[7] Memorandum for the Heads of Executive Offices and Agencies, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," M-24-10 (Mar. 28, 2024), https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf.

[8] ACLU Coalition Letter to OMB Director Russell Vought re. "Executive Order 14179 of January 23, 2025, and Memorandum M-24-10" (Mar. 6, 2025), https://assets.aclu.org/live/uploads/2025/03/03-06-2025-Civil-Society-M-24-10-Letter-aclu-leads-coalition-encouraging-trump-administration-to-preserve-key-ai-safeguardsv3.pdf.

[9] Executive Order 13960 of December 3, 2020, "Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government," 85 Fed. Reg. 78939 (Dec. 8, 2020), https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government; see also Executive Order 13859 of February 11, 2019, "Maintaining American Leadership in Artificial Intelligence," 84 Fed. Reg. 3967 (Feb. 14, 2019)

Memorandum M-24-10 also builds on the first Trump administration's efforts to address potential abuses of and harms from federal uses of AI. M-24-10's minimum risk mitigation practices cover uses of AI that affect Americans' fundamental rights, including uses that restrict protected speech, detect and measure people's thoughts and emotions, surveil children in school, determine terms of employment, and screen applications for mortgages and rental properties.[10] These requirements in M-24-10 not only protect Americans, but also minimize the costs to the government and taxpayers of money wasted on developing harmful AI applications and inviting avoidable litigation.

Transparency and protections against harm provide accountability, which in turn supports public trust and innovation. In the face of growing alarm and distrust about the actions of the Department of Government Efficiency (DOGE),[11] including its efforts to apply AI to Americans' sensitive data,[12] an AI Action Plan should impose clear standards for federal agencies' use of AI and require transparency about public-sector use cases.

Recommendation 5: Invest in developing U.S. open models in order to drive innovation, economic competitiveness, and effective AI diplomacy.
New America's Technology & Democracy Programs have long advocated for the transformative value of open-source software, including open AI models. Last year, New America's Open Technology Institute published a report on the benefits of open models to innovation, security, economic competition, and transparency.[13]

President Trump rightly termed the dramatic emergence of DeepSeek's highly effective open-source models a "wake up call." The essential insight is not that open-source models are dangerous; it is that "[t]he most pressing threat is that China is openly sharing capable AI models that could eventually underpin AI infrastructure across the globe."[14] The country that builds the best and most widely used models will reap benefits for its economy, national security, and global influence. U.S. policymakers should make the diffusion of open-source models developed by U.S. companies a

(recognizing that federal uses of AI must protect "economic and national security, civil liberties, privacy, and American values").
[10] M-24-10, appendix I, sec. 2.
[11] *See, e.g.*, Laurel Wamsley, "The government already knows a lot about you. DOGE is trying to access all of it," NPR (Mar. 11, 2025), https://www.npr.org/2025/03/11/nx-s1-5305054/doge-elon-musk-security-data-information-privacy.
[12] *See, e.g.*, "Hannah Natanson et al., Elon Musk's DOGE is feeding sensitive federal data into AI to target cuts," *The Washington Post* (Feb. 6, 2025), https://www.washingtonpost.com/nation/2025/02/06/elon-musk-doge-ai-department-education/.
[13] Prem Trivedi and Nat Meysenburg, "Openness in Artificial Intelligence Models," *New America* (Nov. 2024), https://www.newamerica.org/oti/reports/openness-in-artificial-intelligence-models/.
[14] Ben Brooks and Michelle Fang, "U.S. leadership in AI requires open-source diplomacy," *The Hill* (Jan. 12, 2025), https://thehill.com/opinion/technology/5079721-china-ai-open-source-threat/ .

cornerstone of our approach to tech diplomacy. Continued U.S. dominance in AI innovation and governance depends upon it.

The AI Action Plan also should recognize that open models provide net benefits to performance and security by inviting broad scrutiny and participation. In addition to allowing researchers to improve models, the ability to examine an AI model by accessing its model weights, training data, and code also allows for formal security vetting by independent third parties. This improves the odds of addressing bias, patching software vulnerabilities, and promoting safety more broadly. These benefits of openness apply both to commercial models and to the intelligence and national security arenas.

At a minimum, U.S. policymakers should avoid restricting open model development, including through the application of heavy-handed export controls that would hamper innovation and hinder U.S. global competitiveness—all without meaningfully restricting peer competitors like China. Instead, the administration should heed the Bipartisan House Task Force Report on AI and the NTIA's recommendations to invest in the open-source AI ecosystem while continuing to gather data for evidence-based assessments of specific marginal risks posed by open models.[15]

Recommendation 6: Forge multi-stakeholder technical partnerships with academia and civil society.

American innovation is at its best when it encourages the private sector to incorporate the expertise and broader perspectives of academia and civil society. The administration should encourage a market-driven approach that leverages American tech sector leadership while incorporating perspectives beyond strictly technical and commercial considerations. The administration should incentivize U.S. companies to establish co-creation frameworks with local academic institutions and civil society organizations that ensure AI systems reflect varied implementation contexts. These partnerships would evaluate technology against both technical excellence and contextual appropriateness, preventing the emergence of technopoly where efficiency metrics alone drive development. This approach maintains U.S. competitive advantage while building systems that are more likely to be widely adopted because they are responsive to local contexts.

[15] Bipartisan House Task Force Report on Artificial Intelligence (Dec. 2024), https://republicans-science.house.gov/_cache/files/a/a/aa2ee12f-8f0c-46a3-8ff8-8e4215d6a72b/A163BDBF496ADA741F831E5BEBBCA06699B6AFF8CC34F4FDC4065BDA298295DF.ai-task-force-report-final.pdf; Dual-Use Foundation Models With Widely Available Model Weights (National Telecommunications and Information Administration, 2024), https://www.ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf.

Recommendation 7: Create a distributed AI monitoring network.
An AI Action plan should prioritize establishing an international consortium to track both technical AI developments and their varying impacts across different economic contexts. This network would analyze emerging capabilities alongside adoption patterns and implementation challenges, providing early warning of both technical risks and problematic power dynamics in AI deployment. The effort should include industry leaders, academic institutions, and civil society organizations from both established and emerging markets to ensure comprehensive insights. This approach provides visibility into technological dependencies that could fuel technopolarization while maintaining U.S. strategic awareness.

Recommendation 8: Catalyze South-South AI Innovation Ecosystems.
The United States has an opportunity to position itself as a strategic enabler of South-South AI exchanges that build technological resilience, reduce dependency risks, and increase U.S. influence globally. The U.S. government should help create lightweight platforms that connect emerging tech hubs across regions with targeted support for knowledge transfer and locally-controlled development resources. A primary focus should be on enabling innovations that address regional priorities through context-appropriate solutions. This approach counters technopolarization by distributing innovation capacity, creates more stable markets for U.S. technologies, and limits the ability of competitor nations to exploit technological dependency narratives.