

**Comments In Response to
The Office of Science and Technology Policy Request for Information on the
Development of an Artificial Intelligence (AI) Action Plan**

March 14, 2025

The U.S. Technology Policy Committee (USTPC) of the Association for Computing Machinery (ACM)¹ is pleased to submit these comments in response to the recent Office of Science and Technology Policy (OSTP) Request for Information (RFI) on the Development of an Artificial Intelligence (AI) Action Plan (“Plan”). The RFI was issued by the Office of Science and Technology Policy on behalf of the Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO) of the National Science Foundation.²

Specifically, USTPC recommends that:

1. AI Definitions and Terms Should Be Harmonized

As noted below, Executive Order 14179,³ upon which the RFI is based, uses the definition of “artificial intelligence” or “AI” as set forth in Chapter 119 of the National Artificial Intelligence Initiative, 15 USC § 9401(3). A shared understanding of core AI terms and concepts is essential as federal agencies, computing professionals, and the broader public develop effective policies for delivering and deploying AI-based services.⁴ Harmonized definitions foster consistency, ensuring everyone works from the same foundation when addressing AI-related challenges. U.S. leadership on the harmonization of definitions and terms in multilateral and other international

¹ The Association for Computing Machinery (ACM), with more than 100,000 members worldwide, is the world’s largest educational and scientific computing society. ACM’s U.S. Technology Policy Committee (USTPC) serves as the focal point for ACM’s interaction with all branches of the U.S. government, the computing community, and the public on policy matters related to information technology.

² Request for Information on the Development of an Artificial Intelligence (AI) Action Plan, 90 FR 9088 (Feb. 6, 2025) <https://www.federalregister.gov/documents/2025/02/06/2025-02305/request-for-information-on-the-development-of-an-artificial-intelligence-ai-action-plan>. This RFI was issued pursuant to Executive Order 14179, issued on January 23, 2025. The Executive Order uses the definition of “artificial intelligence” or “AI” as set forth in [15 U.S.C. 9401\(3\)](https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence):

“The term “artificial intelligence” means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to (A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.”

³ Removing Barriers to American Leadership in Artificial Intelligence, Executive Order 14179, Jan. 23, 2025, <https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence>.

⁴ IEC Editorial Team, “Artificial Intelligence: why terminology matters,” International Electrotechnical Commission, Dec. 15, 2023, <https://www.iec.ch/blog/artificial-intelligence-why-terminology-matters>.

fora will be important to America's ranking in AI innovation and its global market share for AI deployment of AI products and services.

Additionally, agencies should continue emphasizing human-centered design (HCD) when defining and deploying AI systems. HCD is an approach that keeps users' needs, values, and feedback at the core of system design.⁵ It provides methods (like user research, iterative prototyping, and impact assessments) to operationalize these systems consistently. Transparency is improved when AI tools are designed with the end-user in mind with interfaces and explanations that are more understandable to non-experts. HCD encourages clear documentation and reasoning for AI decisions so that users and oversight bodies can trace how outcomes are reached. A critical part of HCD is defining how humans and AI share decision-making.⁶

2. Trust in AI Systems Is Best Achieved Using a Multi-Pronged Approach

Trust is a general term for user acceptance of AI systems. Trust, however, is determined through an analysis of various aspects of an AI system, which may require different metrics (e.g., accuracy, data protection). The accelerating integration of AI into economic, government, and social systems makes resolving distrust of AI not just advisable, but essential, for technological progress. Resistance to AI adoption can cost the global economy in unrealized productivity gains. For example, healthcare systems using AI diagnostics achieve 23% faster treatment times but face 40% slower adoption rates in communities with high algorithmic distrust.⁷

A multi-pronged approach is needed to bridge the trust gap and manage AI risks effectively, including baseline security requirements for foundation model developers and minimum AI security standards for organizations deploying generative AI models (an approach echoed by recent voluntary commitments from leading AI companies.⁸ Equally important are measures to increase transparency about who builds and benefits from AI systems using a task-specific approach, while safeguarding confidentiality and data privacy.⁹

⁵ Artificial Intelligence Risk Management Framework (AI RMF 1.0), National Institute of Standards & Technology, Jan. 2023, <https://doi.org/10.6028/NIST.AI.100-1> at 36.

⁶ "Technology and Empathy: AI in Human-Centered Design, Anthro-Tech, Aug. 29, 2023, <https://anthro-tech.com/ideas-and-insights/ai-impact-on-hcd>.

⁷ The Future of AI-Enabled Health: Leading the Way, World Economic Forum in collaboration with Boston Consulting Group, Jan. 2025, https://reports.weforum.org/docs/WEF_The_Future_of_AI_Enabled_Health_2025.pdf.

⁸ "Voluntary Commitments from Leading Artificial Intelligence Companies on July 21, 2023, 137 Harv. L. Rev. 1284, <https://harvardlawreview.org/print/vol-137/voluntary-commitments-from-leading-artificial-intelligence-companies-on-july-21-2023/>.

⁹ John Zerilli, Umang Bhatt, Adrian Weller, "How transparency modulates trust in artificial intelligence," Patterns (NY), Feb. 24, 2022, doi: [10.1016/j.patter.2022.100455](https://doi.org/10.1016/j.patter.2022.100455).

3. Generative AI Platforms Should be Maintained for Information

Generative AI offers immense benefits but requires safeguards to reduce harm. Providers should maintain public platforms where errors are documented and fixed while alerting users that content may contain inaccuracies. Generative AI applications should disclose their artificial nature, especially when resembling human interactions, and support mechanisms to confirm AI-generated material to avoid confusion or deception.¹⁰

4. AI Governance is the Responsibility of the Implementing and Developing Entities; Policymakers Should Provide Incentives and Initiate Processes for Voluntary Best Practices

The rapid and unpredictable nature of adoption of broadly available AI applications makes governance of the use and deployment of AI by an organization a critical oversight function. Governance of AI should address questions on the responsible party for adopting and using AI in public, institutional, and private settings. AI governance includes an analysis of the AI systems used, the data involved, the algorithmic functions, and use cases to establish policies, processes, and controls to ensure safety, fairness, and compliance.¹¹

The consequences of use must rest with the individual and implementing entity. Governance should address questions on the responsible party for adopting and using AI in public, institutional, and private settings. The reliability of the systems serving as AI may vary significantly without quality control models for applications designed to accomplish critical tasks or provide services/benefits. Vendors should be held to specifications that are enforced through contracts.

Governance may also include oversight of the development of AI applications, the potential for error, or other actions which may impact the user, and how that information is conveyed to the user or corrected. Technical analysis and tools can be used to support the effective governance of AI, which is referred to as technical AI governance (TAIG). TAIG can be used to identify where governance is needed, inform deployment decisions and use of controls, and assess the impacts of a system.¹²

Systems and institutions using algorithmic decision-making should produce explanations regarding the procedures followed and the specific decisions made, especially in public policy contexts. Policymakers should encourage the adoption of mechanisms that enable questioning and redress for those adversely affected by algorithmically-informed decisions. Institutions

¹⁰ Chloe Wittenberg, Ziv Epstein, Adam J. Berinsky, David G. Rand, “Labeling AI-Generated Content: Promises, Perils, and Future Directions,” March 27, 2024, <https://mit-genai.pubpub.org/pub/hu71se89>.

¹¹ See, e.g., “What is AI governance?” IBM, Oct. 10, 2024, <https://www.ibm.com/think/topics/ai-governance>.

¹² Anka Reuel & Ben Bucknall, et al, “Open Problems in Technical AI Governance,” Centre for Governance of AI, July 20, 2024, <https://www.governance.ai/research-paper/open-problems-in-technical-ai-governance>.

should be held responsible by users for the decisions and recommendations made by their algorithms and the results, even if the algorithms' processes are not fully explainable.

Builders of algorithms should maintain descriptions of how training data was collected, including exploration of potential biases induced by data-gathering processes. Testing and training AI systems should be based on real-world use cases from the host organization to develop best practices for how, when, and where its use is justified and valued. Use rigorous methods to validate AI models and document these methods and results, including routine testing to assess and determine whether the model generates discriminatory harm.

Policymakers should provide the incentives to reduce conflict and competition in certain areas to protect the overall health of larger systems that become dependent on AI. Government also may initiate processes in key areas to model best practices for building consensus and voluntary standards to govern the adoption of AI.

5. AI Accountability Should be Enabled By Science-to-Policy Programs and Bug Bounty Programs

Given AI's broad impact, ACM's USTPC has urged researchers to continuously study real-world AI deployments and "periodically make science-based recommendations for their refinement and use."¹³ This mirrors evidence-based approaches in other fields, such as medicine, where expert task forces routinely review new research and update guidelines to improve outcomes. A formal "science-to-policy" pipeline for AI would advance technical research on AI safety, bias, and robustness and inform policies. There are successful analogues in other domains. For instance, in transportation safety, the independent National Transportation Safety Board (NTSB) investigates accidents and then issues safety recommendations which are often adopted.

Bug bounty programs are a proven cybersecurity mechanism that can be adapted to AI systems to bolster their safety and reliability. From machine learning models to AI-driven services, AI systems may have flaws ranging from software bugs to bias or misuse vulnerabilities. Structured bounty programs incentivize independent researchers to stress-test these systems and disclose problems responsibly rather than exploit them.¹⁴

6. Address Challenges to AI Products & Solutions Through Public-Private Initiatives

The rapid deployment of AI products and solutions has created opportunities for fraud, misrepresentations, theft, and other crimes facilitated by AI synthetic data generation capabilities

¹³ Christopher Kang, Jeremy Epstein, Cory Doctorow, Simson Garfinkel, and Jeanna Mathews, "Statement on Principles for the Development and Deployment of Equitable, Private, and Secure Remote Test Administration Systems," ACM, U.S. Technology Policy Committee, Aug. 12, 2021, <https://www.acm.org/binaries/content/assets/public-policy/ustpc-statement-remote-test-admin-systems.pdf>.

¹⁴ Aviram Zrahia, Neil Gandal, Sarit Markovich, Michael Riordan, "The simple economics of an external shock to a bug bounty platform," *Journal of Cybersecurity*, May 8, 2024, <https://academic.oup.com/cybersecurity/article/10/1/tyae006/7667075>.

that can create convincing deep-fakes or illicit content. Similarly, AI hallucinations, where systems generate false information presented as factual, pose subtle yet significant threats to information integrity, especially when hallucinations appear in contexts requiring accuracy, such as legal documents or news sources. To increase and sustain transparency and trust in AI, it is critical to establish mechanisms for content authentication.¹⁵

Countering these challenges to AI products and solutions can be advanced through public-private initiatives that research content authentication, conduct adversarial training on AI models, develop industry-specific guidelines for AI integrity, and establish a clearinghouse for reporting AI-facilitated criminal activities.

7. AI Action Plan Should Include AI Education, Workforce Development, and Research

Artificial Intelligence (AI) education is critically important to maintaining and advancing the United States' position as the global leader in AI. The AI Action Plan should prioritize the following to support innovation and ethical development and deployment of AI:

1. Structured, intentional AI education for all students, evolving from early exposure to AI tools and learning how to use them responsibly and ethically, to understanding algorithmic design, and ultimately to the ability to develop AI solutions and perform research to advance the field.
2. Workforce development that includes continuing AI education, reskilling, and upskilling opportunities to enable workers to contribute meaningfully in a changing industrial and economic landscape to economic growth and national prosperity.
3. Federal research funding and infrastructure that supports collaborations between industry, higher education institutions, and the federal government, to advance basic, translational, and multidisciplinary research in AI and efficient “science-to-policy” and “science-to-product” pipelines.

AI is a rapidly evolving field with transformative potential across all sectors. AI's impact on education, particularly computing education, is profound, changing how we teach, learn, and assess knowledge. Generative AI, a subdiscipline of AI, accelerates this transformation by enabling new ways of learning, communicating, and problem solving. As large language models (LLMs) and AI tools become more powerful, ensuring that all Americans benefit from AI, are educated to use AI effectively, and are prepared to contribute to AI's development in ethical and productive ways is crucial.

8. AI Education Is Essential for AI Leadership and Should Leverage CS2023

Global leadership in AI requires students, researchers, and practitioners to be grounded in a well-rounded, cross-disciplinary education. The quality of AI technology and the applications that

¹⁵ “What are AI hallucinations?” IBM, Sept. 1, 2023, <https://www.ibm.com/think/topics/ai-hallucinations>.

leverage it depend on the quality of the data from which the AI learns and the expertise of those who develop it. In addition, AI is not a standalone discipline; it intersects with computer science, data science, cybersecurity, statistics, engineering, mathematics, and other STEM fields. Applied AI, used across all academic disciplines and industries, requires students to understand the context in which the AI is being used.

In 2023, the ACM, IEEE, and AAAI released CS2023,¹⁶ a three-year effort that resulted in robust computer science curricular guidelines for postsecondary degree-granting programs. These guidelines include AI and are used by institutions worldwide, reflecting the current importance and thought leadership of the United States in computing disciplines. The AI Action Plan should leverage the CS2023 in establishing the US as a leader in AI education.

ACM has provided global thought leadership in computing education and ethics for nearly half a century, making them valuable contributors to AI curricula and policy shaping. Additionally, ACM is a founding member of the Computing Sciences Accreditation Board (CSAB), the lead society for developing and revising criteria for accrediting post-secondary educational programs in various computing disciplines, including AI.

9. The U.S. Should Structure AI Education for Global Competitiveness

To prepare students for a world where AI is increasingly present in all aspects of life, the U.S. should prioritize AI education across five key pathways:

1. Teaching all students AI safety, security, and verification; this is critical for national security.
2. Educating all students on effective AI use, including input optimization and recognizing AI errors.
3. Providing advanced AI development skills to computing students.
4. Incorporating AI applications into various academic disciplines.
5. Workforce development, including reskilling and upskilling workers for AI-driven job transitions.

AI safety and security, sometimes called AI hygiene, is as essential to national security as cybersecurity. Every person in the U.S. who interacts with data or AI must understand what AI is and how to use it without exposing personal or confidential information. Given the proliferation of bad actors using technology for harm, students must learn AI safety and security from a young age.

AI literacy is crucial for AI developers and all students and professionals using AI. Understanding how to structure human interactions with AI, assess AI-generated outputs, and recognize

¹⁶ CS2023: ACM/IEEE-CS/AAAI Computer Science Curricula, ACM, IEEE-CS, AAAI, 2024, <https://csed.acm.org/>.

potential errors or misleading outcomes is critical for effective AI use. AI literacy requires understanding what AI is and is not, and how it works.

As in mathematics, where students progress from basic numeracy to advanced mathematics and statistics, AI education should start with computational thinking—breaking problems into manageable parts—and progress to algorithmic thinking to design and build computer programs and applications. A well-rounded AI curriculum must include technical foundations such as algorithms, machine learning, and data engineering. These foundational skills are essential for all students to support student success, future innovation, and national competitiveness.

Developing national K-12 AI education requirements is challenging due to the decentralized nature of U.S. education, yet high national standards are essential to preparing future AI leaders. Centralized coordination at the federal level is necessary to align AI educational goals across states while allowing flexibility for local implementation.

10. Workforce Development and Reskilling Must be a Priority

Innovation of AI technologies will require a workforce trained to use them. Workforce development in an AI-enhanced world can be divided into three categories:

1. Upskilling and reskilling the general workforce.
2. Continuous learning opportunities for computing practitioners and researchers involved in developing and deploying AI.
3. AI education for teachers.

AI will significantly reshape the American workforce¹⁷, requiring new skills and retraining programs to enable workers to contribute meaningfully and support a robust U.S. economy. Many traditional jobs will evolve, and new roles will emerge, necessitating continuous learning opportunities. The U.S. needs to invest in workforce development initiatives that equip workers with the skills needed for AI-integrated industries.

Successful AI education programs require both advances in teacher education programs for new teachers and ongoing professional development for in-service teachers and faculty. School systems may require additional funding and new certification requirements to attract and retain computer science experts as teachers.

11. AI Education and Research Is Necessary to Sustain America’s AI Leadership

Innovation in AI is driven by research, and a strong education pipeline is necessary to sustain the research enterprise. The U.S. must invest in AI research in universities, national laboratories,

¹⁷ Kristalina Georgieva, “[AI Will Transform the Global Economy. Let’s Make Sure It Benefits Humanity.](https://www.imf.org/en/Blogs/Articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity)” IMF Blog, Jan. 14, 2024, <https://www.imf.org/en/Blogs/Articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity>.

and industry to advance technological advancement. A well-funded research ecosystem will prevent brain drain and sustain America’s AI leadership.

The National Artificial Intelligence Research Resource (NAIRR), a bipartisan and popular entity, plays a critical role in access to AI research infrastructure that neither academia nor private industry can fully provide. Its independence is essential for maintaining neutrality, fostering open innovation, and serving the public interest. By protecting current research investments, the Administration can reinforce trust in NAIRR and maintain access to its capabilities, thereby encouraging broad participation and collaboration in AI innovation. Partnerships with universities and the private sector are proven examples of how technical hubs help establish local expertise. Moreover, NAIRR could invest in alternative AI approaches, especially those designed to minimize resource demands, so that the United States remains at the forefront of sustainable and responsible AI development.

12. Global Competitiveness and AI Education

The United States is a global leader in AI; the IMF AI Preparedness Index places the U.S. in the top three countries globally, along with Singapore and Denmark.¹⁸ The United States benefits from broad academic expertise, a strong startup culture, and a well-funded private sector that fosters innovation. It has robust industry-academia connections, particularly compared to other regions, and leads in chip design and natural resource utilization, which are crucial for energy distribution and cooling. In 2024, the U.S. was home to more than 60% of the world’s top-tier AI researchers.¹⁹ Additionally, the US has a strong foundation in emerging AI domains, giving it a competitive edge over territories such as Europe and Asia. One of its key strengths is its education system, which effectively educates AI development and application talent.

To ensure guarantees of the correctness of AI solutions, deep learning AI system architectures are enhanced with inference-stage, reinforcement learning post-training, and external verification mechanisms. These mechanisms provide critiques or reward signals to adjust the deep learning model’s generation of next candidates towards correct solutions.²⁰ Formal AI verification – including determining whether countries’ AI and AI systems comply with treaty obligations²¹ may require further U.S. investment in education, research and policy. Even as the

¹⁸ Kristalina Georgieva. “[AI Will Transform the Global Economy. Let’s Make Sure It Benefits Humanity.](#)” *IMF Blog*, Jan. 14, 2024.

¹⁹ Jordan Brown. “The Global AI Race: Top 12 Countries Leading the Charge in 2024.” *33 Square*. Updated Jan. 1, 2025, <https://www.33rdsquare.com/top-countries-leading-in-ai-research-technology/>.

²⁰ For insights into verification approaches and novel approaches in DeepSeek R1, see Subbaro Kambhampati, “On the use of Verifiers with LLMs – External vs. Internal LLM-Modulo.” *LinkedIn Blog*, [\(25\) Post | LinkedIn](#).

²¹ Matthew Mittelsteadt, “AI Verification: Mechanisms to Ensure AI Arms Control Compliance,” (Center for Security and Emerging Technology, February 2021). <https://doi.org/10.51593/20190020> or [AI Verification | Center for Security and Emerging Technology](#).

U.S. and China vie to become the world's top AI economy, Canada is the first country with a national AI strategy, and the EU leads the U.S. in guiding the use of AI in democratic societies.²²

It is important to note China's advantage in data access, national subsidies for manufacturing, and AI-driven soft power expansion. As other nations grow in significance in AI education, development, and commercialization, new concerns arise about global dependency on AI solutions developed by countries other than the United States and the ways in which this might weaken U.S. leadership and power.

The U.S. has an opportunity to shape AI education in ways that surpass international efforts, ensuring that American graduates are prepared to continue to lead AI innovation. To maintain its leadership, the U.S. must foster an AI-skilled workforce, support interdisciplinary research, and ensure that AI education is accessible at all levels. AI education will secure America's economic and technological future.

13. Export Controls and National Security Considerations Are Central to U.S. Dominance

Existing export controls²³ on AI technology will be reviewed by the Trump Administration in accordance with the AI Action Plan. Existing controls impose controls on advanced AI chips, cloud access, and the model weights of certain advanced closed-weight (non-open source) dual-use AI models and large clusters of advanced computing integrated circuits (ICs). When China's DeepSeek, a generative AI model that rivaled OpenAI's most advanced model, was released, U.S. tech stock prices fell and downloads of DeepSeek's app exceeded new downloads of ChatGPT.²⁴ Some commentators speculated that U.S. export controls may have provided the impetus for China to figure out a way to perform advanced AI functions without advanced chips.²⁵ A short time later, OpenAI alleged DeepSeek "inappropriately distilled our models."²⁶ In a recent Council of Foreign Relations article, Michael Horowitz notes that the current export control policy does not control open-source models, but if "a good open-source model exists at a certain computational level, controls on the model and weights will no longer succeed."²⁷ Therefore, he

²² Bhaskar Chakravorti, Ajay Bhalla and Ravi Shankar Chaturvedi. "Charting the Emerging Geography of AI." *Harvard Business Review*. Dec. 12, 2023, <https://hbr.org/2023/12/charting-the-emerging-geography-of-ai>.

²³ Framework for Artificial Intelligence Diffusion, U.S. Dept. of Commerce, Bureau of Industry and Security, 90 FR 4544, Jan. 15, 2025, <https://www.govinfo.gov/content/pkg/FR-2025-01-15/pdf/2025-00636.pdf>.

²⁴ "OpenAI says DeepSeek 'inappropriately copied ChatGPT – but it's facing copyright claims too," *The Conversation*, <https://theconversation.com/openai-says-deepseek-inappropriately-copied-chatgpt-but-its-facing-copyright-claims-too-248863> [hereinafter "The Conversation"].

²⁵ Caiwei Chin, "How a top Chinese AI model overcame US sanctions," *MIT Technology Review*, Jan. 24, 2025, <https://www.technologyreview.com/2025/01/24/1110526/china-deepseek-top-ai-despite-sanctions/>.

²⁶ The Conversation.

²⁷ Michael C. Horowitz, "What to Know About the New U.S. AI Diffusion Policy and Export Controls," Council of Foreign Relations, Jan. 13, 2025, <https://www.cfr.org/blog/what-know-about-new-us-ai-diffusion-policy-and-export-controls>.

concludes, “as open-source models improve, more powerful close-weight AI models will be available without restrictions.”²⁸

The U.S. leads in AI innovation, and the largest and most advanced large language model (LLM) databases used for training generative AI and other advanced AI products and services are in the United States. As noted above, export controls that do not restrict access to open-source will provide a gap that will be leveraged by countries subject to the controls. There is a tension between AI export controls on LLMs and global growth by U.S. tech companies. AI is also reshaping global security by enabling adversaries, including non-state actors, to conduct more sophisticated cyberattacks, surveillance, and cyber warfare utilizing autonomous and semi-autonomous weapons. The U.S. and China are in the space race of the 1960s... except it is about AI, not rockets. ACM encourages the Administration to review these national security issues carefully and, in doing so, to leverage the expertise of ACM experts in the areas of AI, export control, and innovation.

14. Diplomatic Efforts for AI Governance Are Important

Diplomatic engagement and inclusive public-private partnerships in multilateral fora, such as G7, OECD, and the United Nations, enables the US to help guide the development of transparent, fair, and secure norms for AI development and advance harmonization of AI in manner that is beneficial to the U.S. Such alliances help prevent fragmented regulations, encourage responsible AI adoption, and uphold democratic values. Engaging industry experts in these efforts helps ensure that emerging policies remain grounded in technical realities.

About This Document:

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.

²⁸ *Id.*