



TECHNOLOGY
POLICY
INSTITUTE

Comments filed with the Office of Science and Technology Policy, Executive Office of the President on “Development of an Artificial Intelligence (AI) Action Plan”

Scott Wallsten, Sarah Oh Lam, and Nathaniel Lovin

March 2025

Before the
Office of Science and Technology Policy
Executive Office of the President
Washington, D.C. 20500

In re:)
Development of an Artificial)
Intelligence (AI) Action Plan) 2025-02305 (90 FR 9088)
)
)

Comments of

Scott Wallsten,¹ Sarah Oh Lam,² and Nathaniel Lovin³

March 14, 2025

Table of Contents

Executive Summary	2
The Fundamentals	2
Define Artificial Intelligence	3
Measure and Monitor First	4
Make Sure Policy Does Not Inhibit Innovation	5
Beware the Precautionary Principle	5
Beware The Temptation to “Pick Winners”	5
Do Not Ignore Potential Harms, But Consider Them Rationally	5
AI R&D Requires The Public and Private Sectors	6
Conclusion	7

¹ Ph.D., President and Senior Fellow, Technology Policy Institute, Washington, D.C.

² J.D., Ph.D., Senior Fellow, Technology Policy Institute, Washington, D.C.

³ Lead Software Engineer and Research Associate, Technology Policy Institute, Washington, D.C.

The views expressed here are those of the authors and do not necessarily reflect those of TPI’s staff, board of directors, or board of academic advisors.

Per filing instructions, the following statement is included: “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.”

Executive Summary

The rapid advancement of artificial intelligence (AI) presents both unprecedented opportunities and challenges. As we strive to maintain America's global leadership in AI innovation, it is crucial that we develop a thoughtful, evidence-based policy framework. This response to the Office of Science and Technology Policy's Request for Information on the Development of an Artificial Intelligence Action Plan outlines key considerations and recommendations for shaping effective AI policy. Our recommendations for priority actions are as follows:

- **Recognize that very little evidence exists about what and even whether to regulate.**
- **Avoid the precautionary principle** that creates rules to preemptively avoid theoretical harms without evidence
- **Resist "picking winners" among competing AI technologies and approaches**, as government is typically ill-placed to make these decisions
- **Involve NIST in developing measurement standards** for AI, leveraging their "global comparative advantage in thinking about measurement"
- **Support government and academic AI research** through both direct federal research at agencies like DARPA, DOD, NASA, and DOE, and robust funding for academic researchers at universities. Both approaches complement private sector innovation and can yield important spillover benefits to the broader economy while addressing areas that may be underinvested by commercial interests
- **Carefully craft and evolve regulatory definitions** of AI to avoid overly broad or narrow rules that might impede innovation
- **Establish systematic evidence collection** about AI benefits and problems rather than relying on anecdotes or hypothetical concerns
- **Develop nuanced risk assessment approaches** for evaluating catastrophic risk scenarios rather than using arbitrary computational power thresholds (as in the Biden EO)
- **Address the fragmented patchwork of state and federal AI regulations** that create compliance challenges
- **Focus on building the knowledge base** necessary for evidence-based policymaking rather than imposing premature constraints

The Fundamentals

At its heart, the way to think about AI policy is no different from any other policy: What outcome do we wish to achieve and how do we achieve it at the lowest cost? That means defining outcomes we want to avoid, behaviors we want to encourage or discourage, and, crucially, understanding the costs of those desires. In other words, as with any other decision, we face the fundamental economic challenge of maximizing net benefits subject to constraints, or alternatively, achieving a desired outcome at the lowest possible cost.

This framework recognizes that policymaking inherently involves tradeoffs. Any regulation that restricts AI development to address potential harms must be evaluated against the

innovation and benefits it might inadvertently prevent. Similarly, any investment to promote AI advancement must consider alternative uses of those resources. Good policy decisions require systematic analysis of these tradeoffs.

A key question in understanding how to reach those objectives is thinking about why the market might not achieve them on its own or why market forces might cause AI to develop in ways society might not prefer. Specifically, it is important to determine whether there are any market failures and externalities not priced into market decisions and, if so, whether policy can address them in a way that increases net benefits. Policymakers might also wish to consider the distribution of benefits and costs across different segments of society alongside efficiency considerations.

Finally, because the future of AI involves significant uncertainty, policymaking in this area requires caution and care. This uncertainty stems from our limited information and experience with the rapidly evolving technology, still-forming market dynamics, and scarce evidence about how various policy approaches might affect AI development. Good policy must therefore incorporate probabilistic thinking about outcomes, considerations of dynamic efficiency—how policy incentives influence innovation trajectories and long-term economic growth—and a focus on building the knowledge base necessary for evidence-based decisions. Thus, many of the recommendations discussed below should be viewed as steps toward developing this essential understanding rather than as definitive judgments.

Define Artificial Intelligence

“Artificial intelligence” has no single accepted definition. The Stanford Artificial Intelligence Laboratory notes that “the lack of a precise, universally accepted definition of AI probably has helped the field to grow, blossom, and advance at an ever-accelerating pace...” while a generally accepted definition by Nils Nilsson is useful: “Artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.”⁴

Policy, however, requires more specificity. Definitions used in rules and regulations are crucial, as they can cause rules to be so broad that they make progress impossible or so narrow that the rules become meaningless. The definition will only become more important as the number of directions in which AI develops and applications in which it is used grows.

The federal government, in 15 U.S.C. § 9401(3), defines “artificial intelligence” as “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;

⁴ Stanford Artificial Intelligence Laboratory, “Defining AI | One Hundred Year Study on Artificial Intelligence (AI100),” (2016), <https://ai100.stanford.edu/2016-report/section-i-what-artificial-intelligence/defining-ai> (last accessed Mar. 13, 2025), citing Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (Cambridge, UK: Cambridge University Press, 2010).

- (B) abstract such perceptions into models through analysis in an automated manner; and
- (C) use model inference to formulate options for information or action.”⁵

This federal definition is not the only one in use. Each state with rules on AI uses its own definition.⁶ Some definitions can be problematic. The National Council of State Legislatures notes that Idaho defines generative AI as “any algorithm or model that creates content such as text, images, audio, or video.”⁷ An algorithm, per se, is not necessarily part of artificial intelligence, and one can imagine this definition being so broad as to affect the use of many types of automation, not just AI.

As AI continues to evolve, regulatory definitions should also evolve and be crafted carefully to avoid ineffective or overly broad rules.

Measure and Monitor First

Clear measures and well-defined outcomes make it possible to collect systematic evidence about benefits and problems rather than relying on anecdotes or hypothetical concerns.

However, arbitrary measurements and thresholds are not likely to be useful. The EO signed by President Biden and California’s SB 1047 set a level of computing power⁸ as the threshold beyond which models should be closely regulated. The threshold level seemed to be arbitrary, and computational power may not be the right metric. Measures should be considered carefully and be targeted at outcomes potentially of concern.

Policymakers can face incentives to pass laws quickly in response to new technologies. AI is no exception. In the last two years, over 1,400 state and federal bills for AI regulations have been introduced, with over 700 bills in 2024 and over 700 bills in 2025 thus far.⁹ In addition to many not being well-reasoned, they create the threat of a fragmented patchwork of compliance requirements for software companies who seek to offer AI services in the United States.

⁵ 15 U.S.C. § 9401(3).

⁶ NCSL, “State Artificial Intelligence (AI) and Related Terms Definition Examples,” <https://www.ncsl.org/technology-and-communication/state-artificial-intelligence-ai-and-related-terms-definition-examples> (last accessed Mar. 13, 2025).

⁷ *Id.*

⁸ Any AI model trained with more than 10^{26} FLOPs (Floating Point operations per Second). *See* Biden AI EO, “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” Oct. 30, 2023, <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>, rescinded by Trump AI EO, Jan. 23, 2025, Removing Barriers to American Leadership in Artificial Intelligence – The White House; California SB 1047, https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB1047.

⁹ Artificial Intelligence (AI) Legislation, MultiState AI, Artificial Intelligence (AI) Legislation, <https://www.multistate.ai/artificial-intelligence-ai-legislation> (last accessed Mar. 7, 2025).

Make Sure Policy Does Not Inhibit Innovation

The public sector should be careful not to impede discovery and experimentation in this pivotal moment of technology acceleration.

Beware the Precautionary Principle

One approach to regulation is the so-called “precautionary principle,” in which policymakers create rules to preemptively avoid possible harms, even without evidence on the likelihood of those harms occurring. Such an approach may be acceptable for an extremely risk-averse population, but can stifle beneficial innovation.

Europe has traditionally embraced the precautionary principle while the U.S. has focused more on addressing harms as we identify them. The EU’s AI Act¹⁰ largely adopts a precautionary principle applied to AI.¹¹ It prohibits certain activities considered “high risk” without evidence that harms would develop while generally requiring developers to prove to the EU that they are safe rather than requiring regulators to identify harms. However, it does attempt to preserve innovation by trying to impose stricter rules on AI that interacts with more critical areas of society and also maintaining “sandboxes” of innovation.

Beware The Temptation to “Pick Winners”

The precautionary principle approach of regulation is not the only way the government can harm innovation. Another is by “picking winners”—deciding which approaches and technology are the right ones and pushing development in that direction. Government is typically ill-placed to make those kinds of decisions and can distort investment and innovation.

An exception is when the government is innovating to improve its own missions and activity. For example, AI research conducted by the Department of Defense, NASA, the Department of Energy, and so on to aid in their own missions can yield important benefits that spill over into the private sector.

Do Not Ignore Potential Harms, But Consider Them Rationally

We do not intend to dismiss potential harms as completely unfounded. AI models can yield suboptimal results, and it is worth thinking hard about what those potential harms could be. In those cases, one key question is whether those outcomes are worse than a typical non-AI outcome and what incentives exist that push developers to correct those biases or not. Another is what the appropriate measures might be to identify those harms, as described above. That makes

¹⁰ EU Artificial Intelligence Act, <https://artificialintelligenceact.eu> (last accessed Mar. 13, 2025).

¹¹ TPI Podcasts: Scott Wallsten and Peter Brown, “AI and Tech in Europe with European Parliament’s Peter Brown” on Two Think Minimum, July 27, 2023, <https://techpolicyinstitute.org/publications/artificial-intelligence/ai-and-tech-in-europe-with-european-parliaments-peter-brown-on-two-think-minimum/>; Nicolas Petit, “European Innovation and Competitiveness with Nicolas Petit,” November 26, 2024, <https://www.twothinkminimum.com/european-innovation-and-competitiveness-with-nicolas-petit/>.

it possible to know objectively whether they occur and helps think through what useful remedies might be.

For example, AI models can yield biased outcomes.¹² One issue in deciding on a response is whether that bias is worse than the bias that would exist otherwise, such as if a human were making the decision instead of an AI. A second issue is that appropriate measurements can be an important step to self-correction—the measurements can highlight the bias and allow developers to address it. Finally, policymakers can consider the problem carefully if problems persist or private actors do not work to correct biases.

A separate class of concerns involves potential existential-level threats from AI systems acting independently of human control. These scenarios are challenging to address coherently as they represent potentially very high-cost but very low-probability events—with reliable estimates for neither. President Biden's AI Executive Order¹³ attempted to regulate such risks, implicitly assuming both high probability and severe consequences without adequately considering the implications of its approach. While the EO correctly identified potential dangers related to AI in chemical, biological, radiological, and nuclear contexts, it relied on an arbitrary computational power threshold as its regulatory trigger rather than a more nuanced risk assessment. The government should certainly consider these catastrophic risk scenarios, but should study them carefully and develop evidence-based approaches rather than implementing precautionary regulations influenced by science fiction narratives that can impede rational discussion.¹⁴

AI R&D Requires The Public and Private Sectors

The private sector is investing tremendous amounts of money in AI in anticipation of extremely large returns.¹⁵ And even though it is impossible to know the socially optimal level of R&D spending, it is hard to believe that there is overall underinvestment that the government needs to address. Indeed, we may come to find out that firms have invested too much.

Nevertheless, the public sector has crucial roles to play. The federal government helped create what we now know as artificial intelligence, with DARPA, for example, spearheading early neural network development.¹⁶ Government continues to advance the field through direct

¹² See, e.g., Cowgill, Bo, and Catherine E. Tucker, “Economics, Fairness and Algorithmic Bias.” *SSRN Electronic Journal*, 2019, <https://doi.org/10.2139/ssrn.3361280>.

¹³ Biden AI EO, “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” Oct. 30, 2023, <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>, rescinded by Trump AI EO, Jan. 23, 2025, “Removing Barriers to American Leadership in Artificial Intelligence – The White House,” <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>.

¹⁴ Some have noted that media discussion of catastrophic risks and their prevalence in popular science fiction have made rational discussion more difficult. See, e.g., Nirit Weiss-Blatt, Adam Thierer, and Taylor Barkley, “The AI Technopanic and Its Effects,” May 2024, <https://abundance.institute/articles/the-ai-technopanic-and-its-effects>.

¹⁵ See, e.g., “Capturing the Growing Artificial Intelligence (AI) Investment Opportunity - iCapital,” <https://icapital.com/insights/private-equity/capturing-the-growing-artificial-intelligence-ai-investment-opportunity/>.

¹⁶ DARPA Neural Network Study: October 1987 - February 1988, <https://s3.amazonaws.com/arena-attachments/2618226/4497a1ae3a5e18d1d9e90a7748c1f9b9.pdf>.

research by its scientists in the military, federal laboratories, and other agencies, and by funding academic researchers working at AI's frontiers.

The government is also crucial in developing the right measures for evaluating AI across fields.¹⁷ This work, while not glamorous, is likely to be a key input into creating frameworks for the government to think about how to create coherent AI policies rather than relying on arbitrary decisions. NIST has a global comparative advantage in thinking about measurement and should be a key part of government decision making.

Conclusion

America's AI leadership depends on taking a measured, evidence-based approach to policy. Rather than rushing to impose restrictions based on hypothetical risks, policymakers should focus on understanding actual outcomes, collecting systematic evidence, and avoiding unnecessary barriers to innovation.

The federal government can play crucial roles in AI development through targeted research funding, measurement standards development, and its own adoption of the technology. However, it should resist both the precautionary principle that has dominated European approaches and the temptation to pick winners among competing technologies.

As we navigate this period of rapid technological change, policy humility is essential. By establishing processes to build knowledge rather than imposing premature constraints, the United States can maintain its competitive edge while developing the expertise needed to address genuine concerns as they emerge with evidence-based solutions.

¹⁷ NIST, "Artificial Intelligence," <https://www.nist.gov/artificial-intelligence> (last accessed Mar. 13, 2025).