MERCATUS CENTER
George Mason University

POLICY BRIEF

# Shaping the AI Revolution: Key Policy Options and Principles for Federal Leaders in 2024

*Matthew Mittelsteadt*

October 2024

When Thomas Edison first demonstrated his electric light bulb, in 1879, the bulbs lasted as little as a few hours or as long as a few days; gas lamps, which Edison was seeking to upend, lasted months. No one deeply understood how electricity worked with any depth; even something as basic as the electron would not be discovered for nearly two decades. Neither electricity generation nor distribution infrastructure existed, save a few research facilities.

By 1930, however, most American factories, offices, and homes were electrified.[1] The change was both rapid and profound. Today society stands on the threshold of a comparatively novel and potentially more potent technological transformation: artificial intelligence (AI).

While most of the adaptation needed to manage AI's many changes will play out naturally, public policy  will play a role in AI's potential revolution, as it has with most modern technical transformations. In coming years, a principal goal of AI policymakers will be striking a balance between safety and innovation. AI has the potential to spur rapid drug development, foster the creation of abundant green energy, and unleash unparalleled progress in manufacturing, among other frontiers. These benefits, however, must not be outweighed by AI's risks, including  AI driven cyberattacks, fraud, misinformation, spam, and other emergent challenges.

Another principal goal for policymakers must be to manage uncertainty. AI's evolution has been rapid and nonstop, and the chief AI policy concerns of the future are unlikely to match those that legislators and regulators try to predict and resolve today.

This policy brief offers a starting point to help federal policymakers navigate AI's many complexities in the coming years, beginning with four broad principles to guide decisions as issues emerge and evolve.

## Principles

AI will create unexpected challenges and necessitate action across many domains. As national leaders formulate legislation and set policy, the following four principles will help guide prudent decisions.

**1. Build State Agility.** In the near term, policymakers should focus on reconfiguring state agencies to adapt rapidly and effectively to the many changes AI could bring. The state should avoid placing too much importance on any one risk without ample evidence. Instead, it should prepare to respond nimbly to variable emergent risks while shoring up its enforcement of existing laws. Building state agility is a function of the following:

- Resources: Federal agencies must have the budgetary, labor, and physical capital needed to implement programs and enforce and implement the law.
- Bureaucratic regulation: Excessive, unclear, complicated, or contradictory rules can needlessly bind decisions and slow action. Managing AI uncertainty will require flexible rules that enable agencies to act and pivot when confronted with unexpected challenges.

**2. Promote Engineering-First Solutions.** Policies will fail if the needed technical tools do not exist. Before considering new regulatory text, decision-makers must first consider whether private sector or state-sponsored engineering fixes can resolve challenges. When regulation is needed, decision-makers must also consider whether regulatory agencies have the necessary technical tools, such as forensic tools and techniques, to implement and enforce the law.

**3. Center AI Diffusion.** A technological invention, no matter how promising, matters little if no one uses it. The 2007 iPhone was a marvel when it debuted, but it did not truly change society until millions of developers could experiment with it and create third-party applications. Diffusion is how a new technology transitions from being a novelty to being an engine of economic productivity.

Widespread use is also how society discovers a technology's limitations, trade-offs, and dangers. Policymakers cannot mitigate AI's downsides without knowing what they are, and they can't acquire that knowledge without allowing extensive practical use of the technology.

Diffusion is an inherently decentralized process, and there are limits to what public policy can encourage. However, the wrong forms of regulation can stymie diffusion. For example, some have proposed that AI developers ensure models are certified free of risk before deployment—an impossibility given the limits of laboratory testing.[2] Regulations of this kind are likely to deter the investment and experimentation needed to arrive at the balance of productivity and safety.

**4. Prioritize Regulating Conduct, Not Models.** Policies with excessive focus on current technology are likely to be rapidly outdated. Society's collective preferences about what should constitute illegal behavior, however, evolve far more slowly. Fraud, assault, theft, and murder have all been considered illicit for centuries. While the means may change, the desire to police such conduct does not.

To remain flexible in the face of rapid change, policymakers should favor regulating conduct, not AI models. Policies should target illicit conduct in a technology-neutral way. Policymakers should also invest in measures that make society less vulnerable to potential threats.

## Federal AI Policy Menu

The options below aim to put these four broad principles into practice. These options can be placed into three categories:

1. **Clarifying regulation** to unleash confident, rapid, responsible AI diffusion.
2. **Investments in state capacity** to enable robust, targeted responses to uncertain AI risks.
3. **Measurement** to spot AI risks and understand AI benefits.

### Clarifying regulation

Institutional uncertainty is a powerful headwind that can work against both AI safety and AI diffusion. Without a clear, consistent, and accessible understanding of the law, private sector innovation and deployment can be chilled. On the safety side, regulatory gaps, overlaps, and contradictions can needlessly bind public and private hands when threats emerge. Policymakers should consider the following options.

#### Option 1: Preemptive regulatory clarity

AI is regulated by several preexisting general-purpose or technology-agnostic statutes. The Food and Drug Administration's (FDA's) medical device approval process, for instance, applies to all devices, whether they include AI or not. As of Fall 2024, the FDA has approved more than 900 such devices.[3] The full list of regulatory authorities that apply to AI, however, remains unclear. Agencies should review statutes and rules and preemptively issue clarifications on their applicability to AI systems.

The Federal Communications Commission (FCC) ruling on the applicability of the Telephone Consumer Protection Act (TCPA), a robocalling regulatory bill, provides a model. The TCPA is a technology-agnostic statute that almost certainly applies to AI-generated audio; by clarifying its applicability, the FCC ensures the industry understands that AI generated audio is out of bounds for robocalling. It should be noted that clarity is distinct from regulatory expansion; limited statutes mustn't be extended beyond their intended reach.

Success depends on implementation. In 2020, the executive order Promoting the Use of Trustworthy Artificial Intelligence in Government required agencies to develop "Agency AI Plans" and proactively identify AI regulatory authorities, yet only 12 percent of agencies complied.[4] This failure was likely due to timing, not substance; President Trump signed the order weeks before the transition to the Biden Administration, and the order did not define who bore responsibility for its implementation. To avoid a retread, the President should assign the Office of Management and Budget (OMB) or another executive agency to coordinate the effort and assign each agency's chief AI officer to implement the requirement. In the event of inaction or further confusion due to a change in administration, Congress should consider legislation to require this effort.[5]

### Option 2: Creation of an AI regulatory map report and tool

AI will interact with many laws and agencies, and the resulting regulatory web risks confusion, overlaps, and contradictions. Legal uncertainty can chill innovation and undermine risk-management efforts such as the National Institute of Standards and Technology (NIST) AI Risk Management Framework.[6]

Once agencies review and clarify laws to assess their AI applicability, Congress should consider commissioning a regulatory map, to be compiled by the OMB or another federal agency with the assistance of responsible agency AI officers. This map would superimpose identified agency authorities and services with existing and likely AI use cases to clarify who would regulate whom. This proposed exercise would have two outputs:

1. Once completed, this map should be assembled into a report for congress,  identifying gaps, overlaps, and contradictions for potential legislative revision.
2. The Department of Commerce should publish the regulatory map as a convenient tool to help the private sector and open-source communities understand how they might be regulated and by whom. This will support the rule of law and enable certain implementation recommendations of the NIST AI Risk Management Framework—notably its recommendation for organizations to understand and review existing law—while allowing the private sector to innovate with greater confidence.

### Option 3: AI-specific regulatory updates

In some cases, existing general-purpose statutes may not fit the unique demands of AI technology. Congress and the executive branch should require agencies to review existing law and determine if statutes are suitable for AI use cases. If an agency finds that a statute cannot be easily applied to AI systems because of the technology's specific characteristics, the agency can recommend adjustments to the law for Congress's consideration. The FDA's proposed regulatory framework for modifications to AI/machine learning-based software as medical devices provides a model. Recognizing that the one-time approval process for medical devices doesn't fit the dynamism of machine learning, the FDA has recommended a more flexible approval process to allow for constant updates while still ensuring device safety.[7]

## State capacity

A second essential tack of near-term federal AI policy is capacity-building. The state needs adequate capabilities and resources to respond to new challenges and to deploy the nonregulatory solutions that will allow for safe, widespread, responsible AI use. In most cases, policymakers should focus on programmatic tune-ups rather than flashy new initiatives. To build capacity, Congress and the executive branch should consider the following options.

### Option 4: Capacity reporting

The state of agency-specific AI capacity is unclear. To assess needs, the executive branch and Congress should task agencies with compiling AI capacity reports that analyze IT gaps, talent needs, and any internal regulations that may stand in the way of agency agility. As a priority, capacity reports should start with technical agencies across branches of government. These include NIST, the OMB, the Cybersecurity and Infrastructure Security Administration (CISA), the Congressional Research Service, the Government Accountability Office, and the Federal Judicial Center. Agencies should share their reports with the OMB and Congress to inform the annual budgeting process.

### Option 5: National vulnerability database investments

AI has already changed the cybersecurity environment. In 2023, large language models yielded a 1,265-percent increase in spear phishing, while machine translation has simultaneously made deceptive attacks more believable. Once deployed, AI systems can be compromised, and the first AI worm was reported in March 2024.[8] In the future, AI could also aid in malware generation, reconnaissance, and autonomous action. Navigating these shifting challenges will require investments in threat-agnostic capacity to shore up existing cybersecurity infrastructure.

NIST's National Vulnerability Database (NVD), devoted to cataloguing, explaining, reporting, and scoring common cyber vulnerabilities, is a key piece of US cybersecurity infrastructure. Engineers use it every day to get up-to-date, reliable information and to scan for threats. Despite its clear utility, the NVD has been experiencing backlogs due to its limited staffing and the increased volume of cyber threats.[9] To keep up with evolving threats and protect the nation's cybersecurity infrastructure, the NVD needs adequate staffing and resources so that new vulnerabilities can be logged without delay.

### Option 6: Open-source threat analysis

CISA estimates that 96 percent of code bases include open-source software (OSS).[10] While OSS comes with the security advantages of crowdsourced analysis, its wide use can carry risks. Recent incidents demonstrate the possibility of systemic risks. The Log4shell vulnerability, often considered the worst in history due to its pervasiveness and the ease of exploitation, impacted 94 percent of cloud environments earlier this decade.

Like all software, AI systems are developed in open-source communities and wield open-source components. AI-generated software threats are likely to follow the example of malicious human actors by taking advantage of commonly used components.

In 2024, CISA released a first-of-its-kind Open-Source Software Security Roadmap that includes the stated intention to "understand OSS software prevalence."[11] Only by understanding the prevalence of software components can decision-makers grasp systemic risks and set priorities for public and private security efforts. This continuous monitoring and analysis of OSS adoption is a sizeable undertaking, yet one that could help protect the most critical systems far into the future. Federal policymakers should invest in the staffing and automated tools needed to guarantee the success of this effort.

### Option 7: Set critical infrastructure priorities

Securing critical infrastructure (CI) against threats is a stated AI priority of the Biden Administration's Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence[12] as well as pending state[13] and federal[14] AI legislation. Lawmakers are working to pass more legislation now and in the future. Therefore, to be sure that their responses are risk-informed, well-targeted, and agile, it's essential that state and federal governments prioritize which CI to protect first in advance of an emergent threat. More than half of the US economy qualifies as CI under current law and policy, including entities like casinos, office buildings, and zoos.[15] To avoid spreading protective efforts too thin, Congress should consider narrowing the definition of CI, either by establishing criteria or by codifying a discrete list of systemically important infrastructure entities or sectors.

### Option 8: AI scholarship for service

Since the signing of the Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, the federal government has received an unprecedented surge in applications for AI-relevant roles.[16] While beneficial, a one-time talent surge cannot guarantee long-run capacity. To ensure a steady flow of AI talent, the federal government should improve talent pipelines.

The CHIPS (Creating Helpful Incentives to Produce Semiconductors) and Science Act authorized an AI Scholarship for Service program to fund post-secondary AI education in exchange for federal service. To build on existing hiring momentum, Congress should immediately authorize the necessary funds.

### Option 9: Forensic tool research and prize challenges

Law enforcement and cybersecurity professionals will require a range of forensic tools to contend with AI's unique dynamics: Generated-content detection applications will be needed to determine the authenticity of evidence, verify the authenticity of content, and identify AI-generated code. Content-provenance tools will be needed to find the source of AI models or the actors behind generated content. When models are fine-tuned versions of general purpose foundation models, model-attribution techniques will be needed to trace and identify those foundation models to inform investigation and supply chain security.[17] Finally, easy to use authenticity techniques and standards will be needed to verify the authenticity of evidence in judicial settings.

Such AI forensic tools should be an immediate priority of federal R&D investments and grants. To diversify investments, federal agencies and Congress should consider prize challenges. Under the America COMPETES Reauthorization Act of 2010, agencies have broad authority to implement and fund innovation prize challenges[18] and should use this authority to crowdsource solutions the federal grant-making process might otherwise miss.

## Measurement and legibility

Effectively prioritizing AI diffusion and state agility depends on measurement, legibility, and understanding. Without a consistent stream of data, problems cannot be identified nor solved. In the AI space, measurement remains a key challenge; existing data is largely incomplete. As a result, problems are misunderstood or blown out of proportion, and solutions are poorly targeted. To improve, the federal government should consider the following options.

### Option 10: Measuring AI diffusion

Already, the US Census Bureau has made modest efforts to take snapshots of AI implementation and diffusion through surveys. Such limited surveys, however, may not capture the full extent of AI diffusion.[19] Understanding the pace and breadth of AI diffusion is essential to grasping AI's influence on labor shifts, key bottlenecks to technical adoption, and potential impacts on economic growth. Solid decisions require solid data. To enable data-informed policymaking, Congress should consider mandating and funding regular, voluntary AI diffusion surveys from the Census Bureau, the Bureau of Labor Statistics, and other relevant statistical agencies.

### Option 11: Ad hoc issue surveys

Accurate threat assessment is also essential to enable targeted, agile responses to challenges. While not everything can be measured, certain problems, like AI-related cybersecurity threats or market instability, have greater salience and are worth understanding in depth. To ensure that the government's understanding of AI challenges remains grounded in reality, NIST's AI Safety Institute should be empowered by Congress to conduct ad hoc issue surveys and studies. Not only would these surveys ground challenges in facts, avoiding the hype common in AI circles, but they would also enable the federal government to take effective action. When problems arise, the federal government can only resolve them quickly if it has identified their existence and put mitigation plans in place.

## Conclusion

AI is evolving at a breakneck pace, as are AI policy, regulation, and legislation. While perhaps modest, the proposals outlined above lean in to present change and uncertainty, working to set a strong institutional foundation for the unknown future. With the right steps today, federal policymakers can set the United States up for success, enabling innovation and resourcing safety efforts, while ensuring rapid, responsible AI diffusion.

## About the Author

Matthew Mittelsteadt is a research fellow and technologist at the Mercatus Center at George Mason University. His work focuses on analyzing AI's impact on national security, cybersecurity, and geopolitics as well as policies that promote rapid, yet responsible diffusion of AI technology in both the private and public sectors. He publishes regularly in his Substack *Digital Spirits*. Prior to joining Mercatus, he worked as a fellow for Syracuse University's Institute for Security Policy and Law. He holds a Master of Science in Cybersecurity from New York University's Tandon School of Engineering and a Master of Public Administration from Syracuse University's Maxwell School.

## Notes

This policy brief is the product of a collaboration between Mercatus scholars Matthew Mittelsteadt and Dean W. Ball.

[1] Jill Jonnes, *Empires of Light: Edison, Tesla, Westinghouse, and the Race to Electrify the World* (Random House Trade, 2004).

[2] Dean W. Ball, "California's Effort to Strangle AI," *Hyperdimensional*, February 9, 2024, https://www.hyperdimensional.co/p/californias-effort-to-strangle-ai.

[3] US Food and Drug Administration, "Artificial Intelligence and Machine Learning (AI/ML)-Enabled Devices," updated August 7, 2024, https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices.

[4] Christine Lawrence, Isaac Cui, and Daniel E. Ho, "Implementation Challenges to Three Pillars of America's AI Strategy" (A white paper for the Stanford Institute for Human-Centered Artificial Intelligence and the Stanford Regulation, Evaluation, and Governance Lab, December 2022), 4.

[5] Legislation seems to have a greater impact than executive orders in this area. Federal agencies also largely ignored the agency AI inventories required under the Trump administration. To remedy this problem, Congress later required the inventories by law, and progress has been made.

[6] Under the "govern" function of the NIST AI Risk Management Framework, the first risk management step organizations must take is to ensure that "[l]egal and regulatory requirements involving AI are understood, managed, and documented." As it is still unclear what laws apply to and regulate AI, it is also unclear as to how this step should be implemented. For groups without robust legal departments, such as the decentralized open-source and small-business communities, this task may be challenging if not impossible. To implement this guidance, regulatory clarity is essential.

[7] US Food and Drug Administration, "Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Service (SaMD): Discussion Paper and Request for Feedback," April 2, 2019.

8. James Coker, “Self-Propagating Worm Created to Target Generative AI Systems,” *Infosecurity Magazine*, March 4, 2024.

9. Edge Editors, “NVD Backlog Continues to Grow,” *Dark Reading*, July 29, 2024.

10. Cybersecurity and Infrastructure Security Agency (CISA), *CISA Open Source Software Security Roadmap*, September 2023.

11. CISA, *Software Security Roadmap.*

12. Exec. Order No. 14110, Fed. Reg. 2023-24283 (October 30, 2023).

13. Safe and Secure Innovation for Frontier Artificial Intelligence Models Act, CA SB-1047 (2023–2024).

14. Office of Sen. John Thune, “Thune, Klobuchar Lead Commerce Committee Colleagues in Introducing Bipartisan AI Bill to Boost Innovation and Strengthen Accountability,” news release, November 15, 2023.

15. Matthew Mittelsteadt, “Critical Risks: Rethinking Critical Infrastructure Policy for Targeted AI Regulation” (Mercatus Policy Brief, Mercatus Center at George Mason University, March 2024).

16. Executive Office of the President of the United States, AI and Tech Talent Task Force, *Increasing AI Capacity Across the Federal Government: AI Talent Surge Progress and Recommendations*, April 2024.

17. Elizabeth Merkhofer, Deepesh Chaudhari, Hyrum S. Anderson, Keith Manville, Lily Wong, and João Gante, “Machine Learning Model Attribution Challenge,” preprint, arXiv, February 13, 2023, https://doi.org/10.48550/arXiv.2302.06716.

18. Marcy A. Gallo, *Federal Prize Competitions* (Congressional Research Service, April 6, 2020).

19. Aakash Kalyani and Marie Hogan, “AI and Productivity Growth: Evidence from Historical Development in Other Technologies,” *On the Economy* (blog), Federal Reserve Bank of St. Louis, April 4, 2024.