PUBLIC SUBMISSION

Received: May 28, 2025
Tracking No. mb8-20gg-3jas
Comments Due: May 28, 2025 **Submission Type:** Web

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0131
Comment on FR Doc # 2025-07332

Submitter Information

Organization: American Society for AI

General Comment

The American Society for AI respectfully submits the attached response to the RFI: "Development of a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan" (Docket ID No. NSF-2025-OGC-0001). Our response outlines strategic recommendations in infrastructure, cybersecurity, ethics, and AI in government. Submitted by Rei Llazani.

Attachments

RFI Response on the Development of a 2025 National AI Research and Development Strategic Plan

American Society for AI

May 28, 2025

Rei Llazani
President
American Society for AI

National Science Foundation
Office of Science and Technology Policy
Networking and Information Technology Research and Development, NCO

RFI Response: "Development of a 2025 National AI R&D Strategic Plan"

Dear NCO, OSTP, and NSF,

We are grateful to the Office of Science and Technology Policy and the National Science Foundation for their leadership in guiding the nation in AI. On behalf of the American Society for AI, we appreciate the opportunity to respond to this RFI.

This response focuses on areas where the U.S. leadership in AI can be strengthened through the development of a strategic plan.

The 2023 Strategic Plan included nine strategies to ensure US leadership in AI.

In this paper, **we propose four (4) additional or overlapping strategies** that must be included in an updated plan:

1. Build Physical Infrastructure (Manufacturing and Energy)
2. Enhance Cybersecurity
3. Cultivate Ethical AI
4. Use AI in Government

Accordingly, in addition to private spending on AI models themselves, the US must invest $2.2 trillion in AI infrastructure over the next five years, including $600 billion for new manufacturing facilities, $1.1 trillion for energy infrastructure, $250 billion for cybersecurity upgrades, $50 billion for safe and ethical AI, and $200 billion for integrating AI in government. Substantive investments in these areas will be critical for ensuring that the US remains competitive, and that AI is used to make the world a better place.

American Society for AI

The American Society for AI (ASFAI) is a private, 501c(3) non-profit organization dedicated to bringing together top AI leaders, executives, & researchers.

Since launching in 2023, **we have emerged as the preeminent AI association** in the United States. Our mission is to make the world a better place with AI, while ensuring the United States remains competitive and *the* leader in AI.

Wholly non-partisan & unbiased, uniquely operating as unpaid volunteers, accepting zero donations & zero sponsorships.

Given our highly esteemed & distinguished membership of senior leaders; including Fortune 100 C-Suite executives, U.S. Senators, and leading AI pioneers, researchers, and entrepreneurs, **we offer unique insights** into both the opportunities and challenges AI presents across industries.

We hope you enjoy our research and report.
We gladly welcome future collaboration and correspondence.

Respectfully submitted on behalf of the Policy Committee of the Society,

Rei Llazani
President
American Society for AI

Key Contributors:

- Distinguished Board Member **Michael Carey**
- Distinguished Member **Russ Wilcox**
- Distinguished Member **Camberley Bates**

Strategy 1: Build Physical Infrastructure

Growth in AI is constrained by the availability of physical infrastructure — specifically, access to high-performance semiconductors and energy resources.

In the short term, advances in AI are most likely to be limited by access to semiconductors (i.e., chips). Advancements in AI have led to explosive growth in the production of advanced chips, and the biggest threat to the availability of these chips is a disruption in global supply chains. The US currently relies heavily on fabrication facilities in Asia for the production of the most advanced semiconductors. Therefore, in the short term the US must do everything possible to ensure that these supply chains are not disrupted. The US must also invest in developing US-based semiconductor manufacturing facilities.

A single semiconductor manufacturing facility for building AI chips can cost $20 billion, and private industry is set to invest over a trillion dollars in such facilities over the next decade. Due to the availability of private funding, streamlining permitting and regulations is more critical than providing direct subsidization for semiconductor manufacturing.

A combination of regulatory and reform subsidization must be adopted to ensure that *at least* $500 billion is invested in US semiconductor manufacturing over the next 5 years.

In the medium term, availability of energy is a more likely bottleneck. Estimates suggest that operating AI systems could consume over 500 TWh of energy annually in the US by 2030 (requiring over 100GW of additional generation capacity). This represents a 15% increase in electricity consumption in the US, not including growth from other areas of the economy, or growth related to manufacturing AI components and other industries that support AI.

To be ready for growth beyond 2030, the US must add approximately 500 GW of generation capacity. Given that electricity generation capacity has been largely stagnant in the US for the last 20 years, it is questionable whether the US energy infrastructure can be made ready for AI without massive intervention.

Thus, government strategy for enhancing the US energy infrastructure must:

1. Enable expansion by streamlining regulatory and permitting processes for a wide range of energy projects
1. Expand short term generation capacity via natural gas generation
2. Build critical national infrastructure for nuclear power
3. Initiate special economic zones (SEZs) targeted toward combined generation capacity and data center location
1. Continue investment in renewable energy
2. Expand, harden, and smarten the energy grid

The US government must act quickly and decisively.
China added over 400 GW of energy production capacity in 2024, compared to about 20 GW in the US. If the US intends to compete with China in the AI space, **1.1 trillion dollars must be invested in energy generation and connection over the next 5 years** (assuming an average cost of $2 billion per GW and an increase of 500 GW of capacity).

Strategy 2: Advance AI Cybersecurity

As AI systems become embedded in critical infrastructure, finance, defense, and healthcare, their vulnerability to cyber threats becomes a national security concern. AI models can be exploited through AI Jailbreaking techniques including, adversarial prompt inputs, data poisoning, model inversion, and API exploitation. And as model weights and training data represent immense economic and strategic value, they have become targets for theft and sabotage. Furthermore, AI models can hallucinate; and as we expand agent-to-agent automation, catastrophic failures are possible. Finally, AI models can be employed to engage in cyberattacks. We have already seen AI used for deepfakes, high-volume scams, highly sophisticated social engineering, and adversarial data collection has increased the overall threat landscape.

The cyber wars have already begun. For example, in 2020, leaked documents revealed that China's Shenzhen Zhenhua Data Technology had compiled a massive database on 2.4 million people worldwide using mostly publicly available information, including tens of thousands of Americans from local officials to military officers. The company systematically harvested public records, social media data, and government databases to create detailed profiles for potential exploitation. Combined with the power of AI, this data could be used to threaten or blackmail American citizens, influence elections, and more.

For example, state and local institutions are exposed in their infrastructure, (utilities, telco), public services (police, fire, education, transportation) where most often cybersecurity is managed at the local level and often lacks the funding or expertise for combating sophisticated AI attacks. Federal support must be given to address these shortcomings. Financial and Healthcare markets contain some of our most sensitive and potentially abusive data if let out in the wild. Successful attacks on these organizations undermine the public trust and ultimately undercut the health of our nation.

Thus, the U.S. must prioritize cybersecurity and security of the models as a foundational pillar of AI development, recognizing that breakthroughs in

performance mean little if the underlying systems are insecure or easily compromised. Ultimately, federal investment is needed across multiple layers of AI cybersecurity to: secure the data supply chain -- ensuring that training data, model weights, and inference endpoints are protected against unauthorized access and tampering; research robust and adversarially resilient AI systems – moving beyond detection and patching toward provable guarantees of model integrity under attack; prepare for the weaponization of AI, both in terms of physical and cyberspace attacks; and develop AI cybersecurity standards and best practices must be advanced – both at the architecture level and the sector level (i.e. State and Local institutions, Finance and Healthcare).

Cybersecurity issues will be as much of a threat to the American way of life as public health concerns. Thus, the US must use the NIH as a model for the scale and methods for investing in cybersecurity. **Accordingly, $250 billion must be allocated to prevent cybersecurity attacks over the next 5 years.**

For example, this could be broken down as follows:

1. AI Counter-Intelligence Capabilities, including offensive and defensive systems, classified programs to disrupt adversarial data collection, attribution systems to identify state-sponsored intelligence gathering, and rapid response capabilities for emerging threats
2. Critical Infrastructure Protection, including hardening of municipal utilities, emergency services, and government continuity systems, integration with critical infrastructure programs, addressing AI-specific vulnerabilities in infrastructure systems, and development of resilience frameworks for cascading failures
3. Classified Research Programs, including AI detection algorithms operating at nation-state scale, quantum-resistant encryption for long-term data protection, adversarial machine learning to predict and counter new attack vectors, and international intelligence sharing platforms with Five Eyes partners
4. Open Research Initiatives, including expansion of NSF Smart and Connected Communities, university partnerships for privacy-preserving technologies, municipal testbeds for real-world validation, and open-source tools for smaller jurisdictions

Cybersecurity must be treated not as an add-on, but as integral to AI design, deployment, and governance. Standards-setting bodies like NIST must expand their AI-specific cybersecurity frameworks, and all government agencies must be required to meet minimum AI security baselines for procurement. Public-private partnerships can accelerate adoption by ensuring that best practices in AI model security are widely disseminated and tested in real-world contexts. Without this level of prioritization, AI's potential could be weaponized by malicious actors faster than it can be safely realized.

Strategy 3: Cultivate Ethical AI

In addition to cybersecurity threats, as the power of AI grows, the potential for harm to the fundamental rights of citizens from the use of AI also grows. For example, AI can harm citizens through negligence, bias, or malicious intent. A national AI strategy must encourage adherence to the principles of responsibility, transparency, and accountability into AI systems.

In a recent survey, ASFAI members made it clear that ensuring AI is used for good is just as important as ensuring rapid advancement in AI capabilities. This includes limitations on how private and public entities use AI. Specifically, some ASFAI members indicated support for:

- Prohibiting autonomous AI agents from engaging in lethal activities without a human in the loop.
- Requiring companies to apply safety measures for AI in high-risk applications, such as applications that use consumer data, or applications involving children.
- Ensuring that generative AI tools can include watermarking mechanisms that embed detectable markers to prevent deepfakes
- Holding companies legally accountable for harms caused by their use of AI
- Ensuring that AI applications avoid algorithmic discrimination and other violations of fundamental rights

Because top-down regulation of AI to ensure responsible growth could come at a heavy cost, care must be taken to ensure that US systems balance the need for growth with the need for ethical application. Specifically, an impact assessment of the EU's AI Act estimated that the cost of preventing violations of fundamental rights could cost 5% of total investment in AI. To ensure that these costs do not slow down the development of AI, the US government must **subsidize the development of safe and ethical AI at a cost of approximately $50 billion over the next 5 years.**

Strategy 4: Use AI in Government

AI has the ability to transform government, to make it more efficient, more transparent, and more democratic. However, this can only happen if governments actually deploy AI to improve government functions. For example, AI systems now possess unprecedented capabilities to extract intelligence from data. This includes the ability to extract and comprehend municipal records that appear routine to human reviewers. As discussed above, this is both a threat and an opportunity. The US must seize the opportunity.

Already, China has announced $1.4 trillion in AI investment over six years — $233 billion annually. This dwarfs total U.S. federal spending across all agencies (~$35 billion annually, comprising approximately $13 billion in civilian agency spending and $14.5 billion in DoD cyber operations). While American cities and states struggle to fund basic IT (80% of cities have fewer than five security personnel), China deploys resources equivalent to the entire U.S. federal.
For example, the Chongqing Big Data Application Development and Management Bureau's 2025 budget allocates 453.91 million RMB ($62.6 million USD) to create comprehensive digital infrastructure for residents of that city.

US cities must be investing at a similar rate, but they cannot solve this challenge individually. The federal government must lead research into AI systems that preserve democratic transparency while preventing weaponization. This requires developing unprecedented technical capabilities including federated AI architectures that enable beneficial municipal collaboration without centralizing sensitive data and privacy-preserving technologies that protect public records from hostile analysis without restricting legitimate transparency.

Industry will not develop these solutions because the municipal market is too fragmented for commercial viability. No company will invest in AI governance frameworks for 30,000 independent jurisdictions with varying technical capabilities, legal requirements, and procurement processes. Federal research investment must address this market failure by developing foundational technologies that municipalities can adapt to their specific needs and constraints.

The research approach must build on existing successful programs while addressing new challenges. NSF's Smart and Connected Communities program currently provides $23.2 million annually through 40 awards supporting interdisciplinary research that integrates intelligent technologies with community challenges. The Department of Homeland Security's State and Local Cybersecurity Grant Program has provided $1 billion over four years (2022-2025) to strengthen local government defenses. NIST's AI Risk Management Framework, released in January 2023, offers voluntary guidance that municipalities can use to assess AI systems.
These investments are not enough.

The US must **invest $200 billion over the next five years to integrate AI in government applications.**

A New American Utility program could be based on successful American infrastructure models like rural electrification.

This could include:

1. A Universal Digital Service Fund, which could be funded by a fee on enterprise cloud computing services, data transfer fees, reallocation of FCC Universal Service Funds, or a dedicated portion of spectrum auction proceeds
2. Regional Digital Infrastructure Cooperatives, including 8-10 regional cooperatives owned by participating municipalities, ember assessments based on population (e.g., $2-3 per resident baseline), service revenues from premium features and commercial applications, and federal loan guarantees modeled on USDA Rural Utilities Service programs
3. State and Local Contributions, including state homeland security grant matching, municipal participation fees scaled by size, regional tax districts for digital infrastructure, public-private partnership investments
4. Private Sector Integration, including mandatory participation for companies above $10B market cap, in-kind contributions of cloud infrastructure and expertise, revenue sharing from commercial applications, and cybersecurity insurance requirements driving investment

This $200 billion annual investment represents approximately a fraction of China's AI spending, and the alternative is accepting technological inferiority and vulnerability across every American community while China builds AI capabilities specifically designed to exploit our open society.

Federal leadership must coordinate these efforts while developing new capabilities specifically for AI-enabled open source governance enhancement. This requires unprecedented partnership between federal laboratories providing fundamental research, universities contributing algorithmic development, and municipal governments serving as real-world testbeds.

This initiative positions America as the global leader in AI systems that enhance rather than undermine democratic governance. International partners will seek American expertise in democratic AI security, creating strategic influence opportunities and technology export markets.

The research outcomes could include: federated AI algorithms, privacy-preserving analysis tools, and distributed governance frameworks, which have applications far beyond municipal government to corporate AI governance, international cooperation, and any system requiring coordination across independent entities.

Conclusion

The U.S. must remain the global leader in AI, and that requires a targeted federal strategy including strengthening critical infrastructure, enhancing security, ensuring ethical AI, and expanding use of AI in government. Doing so will not come without a cost. Specifically, a successful program to ensure US dominance in AI will cost approximately $2.2 trillion over the next 5 years.

Strengthening physical infrastructure, especially domestic chip production and energy supply, is essential to meeting the surging demands of AI development. Similarly, AI-specific cybersecurity must be treated as a national defense priority, with a focus on securing models, data, and systems from both state and non-state threats.

Ethical AI can be advanced without compromising innovation. Instead of imposing rigid rules, the federal government must support best practices through standards, transparency tools, and incentive-based frameworks. Investment in ethical R&D and accountability mechanisms can help ensure AI serves the public good while preserving U.S. competitiveness.

AI must also be used to improve government itself. From federal agencies to small municipalities, strategic funding must empower governments to deploy AI in ways that enhance transparency, automate routine services, and improve access to information. This is especially important for state and local entities, which often lack the resources to adopt advanced technologies.

Taken together -- investments in infrastructure, security, ethics, and government utilization -- form the foundation for long-term AI leadership.