

The WHITE HOUSE

ARTICLES

White House Unveils America's AI Action Plan

The White House

July 23, 2025

The White House today released “Winning the AI Race: America’s AI Action Plan”, in accordance with President Trump’s January executive order on Removing Barriers to American Leadership in AI. Winning the AI race will usher in a new golden age of human flourishing, economic competitiveness, and national security for the American people. The Plan identifies over 90 Federal policy actions across three pillars – Accelerating Innovation, Building American AI Infrastructure, and Leading in International Diplomacy and Security – that the Trump Administration will take in the coming weeks and months. Key policies in the AI Action Plan include:

- **Exporting American AI:** The Commerce and State Departments will partner with industry to deliver secure, full-stack AI export packages – including hardware, models, software, applications, and standards – to America’s friends and allies around the world.
- **Promoting Rapid Buildout of Data Centers:** Expediting and modernizing permits for data centers and semiconductor fabs, as well as creating new national initiatives to increase high-demand occupations like electricians and HVAC technicians.
- **Enabling Innovation and Adoption:** Removing onerous Federal regulations that hinder AI development and deployment, and seek private sector input on rules to remove.
- **Upholding Free Speech in Frontier Models:** Updating Federal procurement guidelines to ensure that the government only contracts with frontier large language model developers who ensure that their systems are objective and free from top-down ideological bias.

“America’s AI Action Plan charts a decisive course to cement U.S. dominance in artificial intelligence. President Trump has prioritized AI as a cornerstone of American innovation, powering a new age of American leadership in science, technology, and global influence. This plan galvanizes Federal efforts to turbocharge our innovation capacity, build cutting-edge infrastructure, and lead globally, ensuring that American workers and families thrive in the AI era. We are moving with urgency to make this vision a reality,” **said White House Office of Science and Technology Policy Director Michael Kratsios.**

“Artificial intelligence is a revolutionary technology with the potential to transform the global economy and alter the balance of power in the world. To remain the leading economic and military power, the United States must win the AI race. Recognizing this, President Trump directed us to produce this Action Plan. To win the AI race, the U.S. must lead in innovation, infrastructure, and global partnerships. At the same time, we must center American workers and avoid Orwellian uses of AI. This Action Plan provides a roadmap for doing that,” **said AI and Crypto Czar David Sacks.**

“Winning the AI Race is non-negotiable. America must continue to be the dominant force in artificial intelligence to promote prosperity and protect our economic and national security. President Trump recognized this at the beginning of his administration and took decisive action by commissioning this AI Action Plan. These clear-cut policy goals set expectations for the Federal Government to ensure America sets the technological gold standard worldwide, and that the world continues to run on American technology,” **said Secretary of State and Acting National Security Advisor Marco Rubio.**

Learn more at [AI.Gov](https://ai.gov).

NEWS

ADMINISTRATION

ISSUES

CONTACT

EOP

VISIT

GALLERY

VIDEO LIBRARY

AMERICA 250

FOUNDING FATHERS



Subscribe to The White House newsletter

Your email

SIGN UP

Text POTUS to 45470 to receive updates

THE WHITE HOUSE

1600 Pennsylvania Ave NW
Washington, DC 20500

WH.GOV

Copyright

Privacy



THE WHITE HOUSE



Winning the Race

AMERICA'S AI ACTION PLAN

JULY 2025

“Today, a new frontier of scientific discovery lies before us, defined by transformative technologies such as artificial intelligence... Breakthroughs in these fields have the potential to reshape the global balance of power, spark entirely new industries, and revolutionize the way we live and work. As our global competitors race to exploit these technologies, it is a national security imperative for the United States to achieve and maintain unquestioned and unchallenged global technological dominance. To secure our future, we must harness the full power of American innovation.”

Donald J. Trump

45th and 47th President of the United States

Table of Contents

Introduction 1

Pillar I: Accelerate AI Innovation 3

- Remove Red Tape and Onerous Regulation 3
- Ensure that Frontier AI Protects Free Speech and American Values 4
- Encourage Open-Source and Open-Weight AI 4
- Enable AI Adoption 5
- Empower American Workers in the Age of AI 6
- Support Next-Generation Manufacturing 7
- Invest in AI-Enabled Science 8
- Build World-Class Scientific Datasets 8
- Advance the Science of AI 9
- Invest in AI Interpretability, Control, and Robustness Breakthroughs 9
- Build an AI Evaluations Ecosystem 10
- Accelerate AI Adoption in Government 10
- Drive Adoption of AI within the Department of Defense 11
- Protect Commercial and Government AI Innovations 12
- Combat Synthetic Media in the Legal System 12

Pillar II: Build American AI Infrastructure 14

- Create Streamlined Permitting for Data Centers, Semiconductor Manufacturing Facilities, and Energy Infrastructure while Guaranteeing Security 14
- Develop a Grid to Match the Pace of AI Innovation 15
- Restore American Semiconductor Manufacturing 16
- Build High-Security Data Centers for Military and Intelligence Community Usage 16
- Train a Skilled Workforce for AI Infrastructure 17
- Bolster Critical Infrastructure Cybersecurity 18
- Promote Secure-By-Design AI Technologies and Applications 18
- Promote Mature Federal Capacity for AI Incident Response 19

Pillar III: Lead in International AI Diplomacy and Security 20

- Export American AI to Allies and Partners 20
- Counter Chinese Influence in International Governance Bodies 20
- Strengthen AI Compute Export Control Enforcement 21
- Plug Loopholes in Existing Semiconductor Manufacturing Export Controls 21
- Align Protection Measures Globally 21
- Ensure that the U.S. Government is at the Forefront of Evaluating National Security Risks in Frontier Models 22
- Invest in Biosecurity 23

Introduction

The United States is in a race to achieve global dominance in artificial intelligence (AI). Whoever has the largest AI ecosystem will set global AI standards and reap broad economic and military benefits. Just like we won the space race, it is imperative that the United States and its allies win this race. President Trump took decisive steps toward achieving this goal during his first days in office by signing Executive Order 14179, “Removing Barriers to American Leadership in Artificial Intelligence,” calling for America to retain dominance in this global race and directing the creation of an AI Action Plan.¹

Winning the AI race will usher in a new golden age of human flourishing, economic competitiveness, and national security for the American people. AI will enable Americans to discover new materials, synthesize new chemicals, manufacture new drugs, and develop new methods to harness energy—an industrial revolution. It will enable radically new forms of education, media, and communication—an information revolution. And it will enable altogether new intellectual achievements: unraveling ancient scrolls once thought unreadable, making breakthroughs in scientific and mathematical theory, and creating new kinds of digital and physical art—a *renaissance*.

An industrial revolution, an information revolution, and a renaissance—all at once. This is the potential that AI presents. The opportunity that stands before us is both inspiring and humbling. And it is ours to seize, or to lose.

America’s AI Action Plan has three pillars: innovation, infrastructure, and international diplomacy and security. The United States needs to innovate faster and more comprehensively than our competitors in the development and distribution of new AI technology across every field, and dismantle unnecessary regulatory barriers that hinder the private sector in doing so. As Vice President Vance remarked at the Paris AI Action Summit in February, restricting AI development with onerous regulation “would not only unfairly benefit incumbents... it would mean paralyzing one of the most promising technologies we have seen in generations.”² That is why President Trump rescinded the Biden Administration’s dangerous actions on day one.

We need to build and maintain vast AI infrastructure and the energy to power it. To do that, we will continue to reject radical climate dogma and bureaucratic red tape, as the Administration has done since Inauguration Day. Simply put, we need to “Build, Baby, Build!”

We need to establish American AI—from our advanced semiconductors to our models to our applications—as the gold standard for AI worldwide and ensure our allies are building on American technology.

Several principles cut across each of these three pillars. First, American workers are central to the Trump Administration’s AI policy. The Administration will ensure that our Nation’s workers and their families gain from the opportunities created in this technological revolution. The AI infrastructure buildout will create high-paying jobs for American workers. And the

¹ Executive Order 14179 of January 23, 2025, “Removing Barriers to American Leadership in Artificial Intelligence,” Federal Register 90 (20) 8741, www.govinfo.gov/content/pkg/FR-2025-01-31/pdf/2025-02172.pdf.

² J.D. Vance, “Remarks by the Vice President at the Artificial Intelligence Action Summit in Paris, France,” February 11, 2025, www.presidency.ucsb.edu/documents/remarks-the-vice-president-the-artificial-intelligence-action-summit-paris-france.

breakthroughs in medicine, manufacturing, and many other fields that AI will make possible will increase the standard of living for all Americans. AI will improve the lives of Americans by complementing their work—not replacing it.

Second, our AI systems must be free from ideological bias and be designed to pursue objective truth rather than social engineering agendas when users seek factual information or analysis. AI systems are becoming essential tools, profoundly shaping how Americans consume information, but these tools must also be trustworthy.

Finally, we must prevent our advanced technologies from being misused or stolen by malicious actors as well as monitor for emerging and unforeseen risks from AI. Doing so will require constant vigilance.

This Action Plan sets forth clear policy goals for near-term execution by the Federal government. The Action Plan's objective is to articulate policy recommendations that this Administration can deliver for the American people to achieve the President's vision of global AI dominance. The AI race is America's to win, and this Action Plan is our roadmap to victory.

Michael J. Kratsios
Assistant to the President for Science and Technology

David O. Sacks
Special Advisor for AI and Crypto

Marco A. Rubio
Assistant to the President for National Security Affairs

Pillar I: Accelerate AI Innovation

America must have the most powerful AI systems in the world, but we must also lead the world in creative and transformative application of these systems. Achieving these goals requires the Federal government to create the conditions where private-sector-led innovation can flourish.

Remove Red Tape and Onerous Regulation

To maintain global leadership in AI, America's private sector must be unencumbered by bureaucratic red tape. President Trump has already taken multiple steps toward this goal, including rescinding Biden Executive Order 14110 on AI that foreshadowed an onerous regulatory regime.³ AI is far too important to smother in bureaucracy at this early stage, whether at the state or Federal level. The Federal government should not allow AI-related Federal funding to be directed toward states with burdensome AI regulations that waste these funds, but should also not interfere with states' rights to pass prudent laws that are not unduly restrictive to innovation.

Recommended Policy Actions

- Led by the Office of Science and Technology Policy (OSTP), launch a Request for Information from businesses and the public at large about current Federal regulations that hinder AI innovation and adoption, and work with relevant Federal agencies to take appropriate action.
- Led by the Office of Management and Budget (OMB) and consistent with Executive Order 14192 of January 31, 2025, "Unleashing Prosperity Through Deregulation," work with all Federal agencies to identify, revise, or repeal regulations, rules, memoranda, administrative orders, guidance documents, policy statements, and interagency agreements that unnecessarily hinder AI development or deployment.⁴
- Led by OMB, work with Federal agencies that have AI-related discretionary funding programs to ensure, consistent with applicable law, that they consider a state's AI regulatory climate when making funding decisions and limit funding if the state's AI regulatory regimes may hinder the effectiveness of that funding or award.
- Led by the Federal Communications Commission (FCC), evaluate whether state AI regulations interfere with the agency's ability to carry out its obligations and authorities under the Communications Act of 1934.⁵
- Review all Federal Trade Commission (FTC) investigations commenced under the previous administration to ensure that they do not advance theories of liability that unduly burden AI innovation. Furthermore, review all FTC final orders, consent decrees,

³ Executive Order 14110 of October 30, 2023, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," Federal Register 88 (210) 75191, www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf.

⁴ Executive Order 14192 of January 31, 2025, "Unleashing Prosperity Through Deregulation," Federal Register 90 (24) 9065, www.govinfo.gov/content/pkg/FR-2025-02-06/pdf/2025-02345.pdf.

⁵ Communications Act of 1934, 47 U.S.C. §§ 151-646.

and injunctions, and, where appropriate, seek to modify or set-aside any that unduly burden AI innovation.

Ensure that Frontier AI Protects Free Speech and American Values

AI systems will play a profound role in how we educate our children, do our jobs, and consume media. It is essential that these systems be built from the ground up with freedom of speech and expression in mind, and that U.S. government policy does not interfere with that objective. We must ensure that free speech flourishes in the era of AI and that AI procured by the Federal government objectively reflects truth rather than social engineering agendas.

Recommended Policy Actions

- Led by the Department of Commerce (DOC) through the National Institute of Standards and Technology (NIST), revise the NIST AI Risk Management Framework to eliminate references to misinformation, Diversity, Equity, and Inclusion, and climate change.⁶
- Update Federal procurement guidelines to ensure that the government only contracts with frontier large language model (LLM) developers who ensure that their systems are objective and free from top-down ideological bias.
- Led by DOC through NIST's Center for AI Standards and Innovation (CAISI), conduct research and, as appropriate, publish evaluations of frontier models from the People's Republic of China for alignment with Chinese Communist Party talking points and censorship.

Encourage Open-Source and Open-Weight AI

Open-source and open-weight AI models are made freely available by developers for anyone in the world to download and modify. Models distributed this way have unique value for innovation because startups can use them flexibly without being dependent on a closed model provider. They also benefit commercial and government adoption of AI because many businesses and governments have sensitive data that they cannot send to closed model vendors. And they are essential for academic research, which often relies on access to the weights and training data of a model to perform scientifically rigorous experiments.

We need to ensure America has leading open models founded on American values. Open-source and open-weight models could become global standards in some areas of business and in academic research worldwide. For that reason, they also have geostrategic value. While the decision of whether and how to release an open or closed model is fundamentally up to the developer, the Federal government should create a supportive environment for open models.

Recommended Policy Actions

- Ensure access to large-scale computing power for startups and academics by improving the financial market for compute. Currently, a company seeking to use large-scale compute must often sign long-term contracts with hyperscalers—far beyond the

⁶ National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," (Gaithersburg, MD: National Institute of Standards and Technology, 2023), www.doi.org/10.6028/NIST.AI.100-1.

budgetary reach of most academics and many startups. America has solved this problem before with other goods through financial markets, such as spot and forward markets for commodities. Through collaboration with industry, NIST at DOC, OSTP, and the National Science Foundation's (NSF) National AI Research Resource (NAIRR) pilot, the Federal government can accelerate the maturation of a healthy financial market for compute.

- Partner with leading technology companies to increase the research community's access to world-class private sector computing, models, data, and software resources as part of the NAIRR pilot.
- Build the foundations for a lean and sustainable NAIRR operations capability that can connect an increasing number of researchers and educators across the country to critical AI resources.
- Continue to foster the next generation of AI breakthroughs by publishing a new National AI Research and Development (R&D) Strategic Plan, led by OSTP, to guide Federal AI research investments.
- Led by DOC through the National Telecommunications and Information Administration (NTIA), convene stakeholders to help drive adoption of open-source and open-weight models by small and medium-sized businesses.

Enable AI Adoption

Today, the bottleneck to harnessing AI's full potential is not necessarily the availability of models, tools, or applications. Rather, it is the limited and slow adoption of AI, particularly within large, established organizations. Many of America's most critical sectors, such as healthcare, are especially slow to adopt due to a variety of factors, including distrust or lack of understanding of the technology, a complex regulatory landscape, and a lack of clear governance and risk mitigation standards. A coordinated Federal effort would be beneficial in establishing a dynamic, "try-first" culture for AI across American industry.

Recommended Policy Actions

- Establish regulatory sandboxes or AI Centers of Excellence around the country where researchers, startups, and established enterprises can rapidly deploy and test AI tools while committing to open sharing of data and results. These efforts would be enabled by regulatory agencies such as the Food and Drug Administration (FDA) and the Securities and Exchange Commission (SEC), with support from DOC through its AI evaluation initiatives at NIST.
- Launch several domain-specific efforts (e.g., in healthcare, energy, and agriculture), led by NIST at DOC, to convene a broad range of public, private, and academic stakeholders to accelerate the development and adoption of national standards for AI systems and to measure how much AI increases productivity at realistic tasks in those domains.
- Led by the Department of Defense (DOD) in coordination with the Office of the Director of National Intelligence (ODNI), regularly update joint DOD-Intelligence Community (IC) assessments of the comparative level of adoption of AI tools by the United States, its competitors, and its adversaries' national security establishments, and establish an

approach for continuous adaptation of the DOD and IC's respective AI adoption initiatives based on these AI net assessments.

- Prioritize, collect, and distribute intelligence on foreign frontier AI projects that may have national security implications, via collaboration between the IC, the Department of Energy (DOE), CAISI at DOC, the National Security Council (NSC), and OSTP.

Empower American Workers in the Age of AI

The Trump Administration supports a worker-first AI agenda. By accelerating productivity and creating entirely new industries, AI can help America build an economy that delivers more pathways to economic opportunity for American workers. But it will also transform how work gets done across all industries and occupations, demanding a serious workforce response to help workers navigate that transition. The Trump Administration has already taken significant steps to lead on this front, including the April 2025 Executive Orders 14277 and 14278, “Advancing Artificial Intelligence Education for American Youth” and “Preparing Americans for High-Paying Skilled Trade Jobs of the Future.”^{7,8} To continue delivering on this vision, the Trump Administration will advance a priority set of actions to expand AI literacy and skills development, continuously evaluate AI's impact on the labor market, and pilot new innovations to rapidly retrain and help workers thrive in an AI-driven economy.

Recommended Policy Actions

- Led by the Department of Labor (DOL), the Department of Education (ED), NSF, and DOC, prioritize AI skill development as a core objective of relevant education and workforce funding streams. This should include promoting the integration of AI skill development into relevant programs, including career and technical education (CTE), workforce training, apprenticeships, and other federally supported skills initiatives.
- Led by the Department of the Treasury, issue guidance clarifying that many AI literacy and AI skill development programs may qualify as eligible educational assistance under Section 132 of the Internal Revenue Code, given AI's widespread impact reshaping the tasks and skills required across industries and occupations.⁹ In applicable situations, this will enable employers to offer tax-free reimbursement for AI-related training and help scale private-sector investment in AI skill development, preserving jobs for American workers.
- Led by the Bureau of Labor Statistics (BLS) and DOC through the Census Bureau and the Bureau of Economic Analysis (BEA), study AI's impact on the labor market by using data they already collect on these topics, such as the firm-level AI adoption trends the Census Bureau tracks in its Business Trends and Outlook Survey. These agencies could then provide analysis of AI adoption, job creation, displacement, and wage effects.
- Establish the AI Workforce Research Hub under DOL to lead a sustained Federal effort to evaluate the impact of AI on the labor market and the experience of the American

⁷ Executive Order 14277 of April 23, 2025: “Advancing Artificial Intelligence Education for American Youth,” Federal Register 90 (80) 17519, www.govinfo.gov/content/pkg/FR-2025-04-28/pdf/2025-07368.pdf.

⁸ Executive Order 14278 of April 23, 2025: “Preparing Americans for High-Paying Skilled Trade Jobs of the Future,” Federal Register 90 (80) 17525, www.govinfo.gov/content/pkg/FR-2025-04-28/pdf/2025-07369.pdf.

⁹ Revenue Act of 1978, 26 U.S.C. § 132.

worker, in collaboration with BLS and DOC through the Census Bureau and BEA. The Hub would produce recurring analyses, conduct scenario planning for a range of potential AI impact levels, and generate actionable insights to inform workforce and education policy.

- Led by DOL, leverage available discretionary funding, where appropriate, to fund rapid retraining for individuals impacted by AI-related job displacement. Issue clarifying guidance to help states identify eligible dislocated workers in sectors undergoing significant structural change tied to AI adoption, as well as guidance clarifying how state Rapid Response funds can be used to proactively upskill workers at risk of future displacement.
- At DOL and DOC, rapidly pilot new approaches to workforce challenges created by AI, which may include areas such as rapid retraining needs driven by worker displacement and shifting skill requirements for entry-level roles. These pilots should be carried out by states and workforce intermediaries using existing authorities under the Workforce Innovation and Opportunity Act and the Public Works and Economic Development Act, and should be designed to identify surface scalable, performance-driven strategies that help the workforce system adapt to the speed and complexity of AI-driven labor market change.^{10, 11}

Support Next-Generation Manufacturing

AI will enable a wide range of new innovations in the physical world: autonomous drones, self-driving cars, robotics, and other inventions for which terminology does not yet exist. It is crucial that America and our trusted allies be world-class manufacturers of these next-generation technologies. AI, robotics, and related technologies create opportunities for novel capabilities in manufacturing and logistics, including ones with applications to defense and national security. The Federal government should prioritize investment in these emerging technologies and usher in a new industrial renaissance.

Recommended Policy Actions

- Invest in developing and scaling foundational and translational manufacturing technologies via DOD, DOC, DOE, NSF, and other Federal agencies using the Small Business Innovation Research program, the Small Business Technology Transfer program, research grants, CHIPS R&D programs, Stevenson-Wydler Technology Innovation Act authorities, Title III of the Defense Production Act, Other Transaction Authority, and other authorities.^{12, 13, 14, 15}
- Led by DOC through NTIA, convene industry and government stakeholders to identify supply chain challenges to American robotics and drone manufacturing.

¹⁰ Workforce Innovation and Opportunity Act of 2014, 29 U.S.C. §§ 3101-3361.

¹¹ Public Works and Economic Development Act of 1965, 42 U.S.C. §§ 3121-3233.

¹² William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 15 U.S.C. § 4656.

¹³ Stevenson-Wydler Technology Innovation Act of 1980, Pub. L. No. 96-480, 94 Stat. 2311 (codified as amended in scattered sections of 15 U.S.C.).

¹⁴ Defense Production Act of 1950, 50 U.S.C. §§ 4551-4568.

¹⁵ National Defense Authorization Act for Fiscal years 1990 and 1991, 10 U.S.C. §§ 4021-4022.

Invest in AI-Enabled Science

Like many other domains, science itself will be transformed by AI. AI systems can already generate models of protein structures, novel materials, and much else. Increasingly powerful general-purpose models show promise in formulating hypotheses and designing experiments. These nascent capabilities promise to accelerate scientific advancement. They will only do so, however, with critical changes in the way science is conducted, including the enabling scientific infrastructure. AI-enabled predictions are of little use if scientists cannot also increase the scale of experimentation. Basic science today is often a labor-intensive process; the AI era will require more scientific and engineering research to transform theories into industrial-scale enterprises. This, in turn, will necessitate new infrastructure and support of new kinds of scientific organizations.

Recommended Policy Actions

- Through NSF, DOE, NIST at DOC, and other Federal partners, invest in automated cloud-enabled labs for a range of scientific fields, including engineering, materials science, chemistry, biology, and neuroscience, built by, as appropriate, the private sector, Federal agencies, and research institutions in coordination and collaboration with DOE National Laboratories.
- Use long-term agreements to support Focused-Research Organizations or other similar entities using AI and other emerging technologies to make fundamental scientific advancements.
- Incentivize researchers to release more high-quality datasets publicly by considering the impact of scientific and engineering datasets from a researchers' prior funded efforts in the review of proposals for new projects.
- Require federally funded researchers to disclose non-proprietary, non-sensitive datasets that are used by AI models during the course of research and experimentation.

Build World-Class Scientific Datasets

High-quality data has become a national strategic asset as governments pursue AI innovation goals and capitalize on the technology's economic benefits. Other countries, including our adversaries, have raced ahead of us in amassing vast troves of scientific data. The United States must lead the creation of the world's largest and highest quality AI-ready scientific datasets, while maintaining respect for individual rights and ensuring civil liberties, privacy, and confidentiality protections.

Recommended Policy Actions

- Direct the National Science and Technology Council (NSTC) Machine Learning and AI Subcommittee to make recommendations on minimum data quality standards for the use of biological, materials science, chemical, physical, and other scientific data modalities in AI model training.
- Promulgate the OMB regulations required in the Confidential Information Protection and Statistical Efficiency Act of 2018 on presumption of accessibility and expanding secure access, which will lower barriers and break down silos to accessing Federal data,

ultimately facilitating the improved use of AI for evidence building by statistical agencies while protecting confidential data from inappropriate access and use.¹⁶

- Establish secure compute environments within NSF and DOE to enable secure AI use-cases for controlled access to restricted Federal data.
- Create an online portal for NSF's National Secure Data Service (NSDS) demonstration project to provide the public and Federal agencies with a front door to AI use-cases involving controlled access to restricted Federal data.
- Explore the creation of a whole-genome sequencing program for life on Federal lands, led by the NSTC and including members of the U.S. Department of Agriculture, DOE, NIH, NSF, the Department of Interior, and Cooperative Ecosystem Studies Units to collaborate on the development of an initiative to establish a whole genome sequencing program for life on Federal lands (to include all biological domains). This new data would be a valuable resource in training future biological foundation models.

Advance the Science of AI

Just as LLMs and generative AI systems represented a paradigm shift in the science of AI, future breakthroughs may similarly transform what is possible with AI. It is imperative that the United States remain the leading pioneer of such breakthroughs, and this begins with strategic, targeted investment in the most promising paths at the frontier.

Recommended Policy Actions

- Prioritize investment in theoretical, computational, and experimental research to preserve America's leadership in discovering new and transformative paradigms that advance the capabilities of AI, reflecting this priority in the forthcoming National AI R&D Strategic Plan.

Invest in AI Interpretability, Control, and Robustness Breakthroughs

Today, the inner workings of frontier AI systems are poorly understood. Technologists know how LLMs work at a high level, but often cannot explain why a model produced a specific output. This can make it hard to predict the behavior of any specific AI system. This lack of predictability, in turn, can make it challenging to use advanced AI in defense, national security, or other applications where lives are at stake. The United States will be better able to use AI systems to their fullest potential in high-stakes national security domains if we make fundamental breakthroughs on these research problems.

Recommended Policy Actions

- Launch a technology development program led by the Defense Advanced Research Projects Agency in collaboration with CAISI at DOC and NSF, to advance AI interpretability, AI control systems, and adversarial robustness.

¹⁶ Confidential Information Protection and Statistical Efficiency Act of 2018, 44 U.S.C. §§ 3561-3583.

- Prioritize fundamental advancements in AI interpretability, control, and robustness as part of the forthcoming National AI R&D Strategic Plan.
- The DOD, DOE, CAISI at DOC, the Department of Homeland Security (DHS), NSF, and academic partners should coordinate an AI hackathon initiative to solicit the best and brightest from U.S. academia to test AI systems for transparency, effectiveness, use control, and security vulnerabilities.

Build an AI Evaluations Ecosystem

Evaluations are how the AI industry assesses the performance and reliability of AI systems. Rigorous evaluations can be a critical tool in defining and measuring AI reliability and performance in regulated industries. Over time, regulators should explore the use of evaluations in their application of existing law to AI systems.

Recommended Policy Actions

- Publish guidelines and resources through NIST at DOC, including CAISI, for Federal agencies to conduct their own evaluations of AI systems for their distinct missions and operations and for compliance with existing law.
- Support the development of the science of measuring and evaluating AI models, led by NIST at DOC, DOE, NSF, and other Federal science agencies.
- Convene meetings at least twice per year under the auspices of CAISI at DOC for Federal agencies and the research community to share learnings and best practices on building AI evaluations.
- Invest, via DOE and NSF, in the development of AI testbeds for piloting AI systems in secure, real-world settings, allowing researchers to prototype new AI systems and translate them to the market. Such testbeds would encourage participation by broad multistakeholder teams and span a wide variety of economic verticals touched by AI, including agriculture, transportation, and healthcare delivery.
- Led by DOC, convene the NIST AI Consortium to empower the collaborative establishment of new measurement science that will enable the identification of proven, scalable, and interoperable techniques and metrics to promote the development of AI.

Accelerate AI Adoption in Government

With AI tools in use, the Federal government can serve the public with far greater efficiency and effectiveness. Use cases include accelerating slow and often manual internal processes, streamlining public interactions, and many others. Taken together, transformative use of AI can help deliver the highly responsive government the American people expect and deserve.

OMB has already advanced AI adoption in government by reducing onerous rules imposed by the Biden Administration.^{17, 18} Now is the time to build on this success.

Recommended Policy Actions

- Formalize the Chief Artificial Intelligence Officer Council (CAIOC) as the primary venue for interagency coordination and collaboration on AI adoption. Through the CAIOC, initiate strategic coordination and collaboration with relevant Federal executive councils, to include: the President's Management Council, Chief Data Officer Council, Chief Information Officer Council, Interagency Council on Statistical Policy, Chief Human Capital Officer Council, and Federal Privacy Council.
- Create a talent-exchange program designed to allow rapid details of Federal staff to other agencies in need of specialized AI talent (e.g., data scientists and software engineers), with input from the Office of Personnel Management.
- Create an AI procurement toolbox managed by the General Services Administration (GSA), in coordination with OMB, that facilitates uniformity across the Federal enterprise to the greatest extent practicable. This system would allow any Federal agency to easily choose among multiple models in a manner compliant with relevant privacy, data governance, and transparency laws. Agencies should also have ample flexibility to customize models to their own ends, as well as to see a catalog of other agency AI uses (based on OMB's pre-existing AI Use Case Inventory).
- Implement an Advanced Technology Transfer and Capability Sharing Program with GSA to quickly transfer advanced AI capabilities and use cases between agencies.
- Mandate that all Federal agencies ensure—to the maximum extent practicable—that all employees whose work could benefit from access to frontier language models have access to, and appropriate training for, such tools.
- Convene, under the auspices of OMB, a cohort of agencies with High Impact Service Providers to pilot and increase the use of AI to improve the delivery of services to the public.

Drive Adoption of AI within the Department of Defense

AI has the potential to transform both the warfighting and back-office operations of the DOD. The United States must aggressively adopt AI within its Armed Forces if it is to maintain its global military preeminence while also ensuring, as outlined throughout this Action Plan, that its use of AI is secure and reliable. Because the DOD has unique operational needs within the Federal government, it merits specific policy actions to drive AI adoption.

¹⁷ Office of Management and Budget, "Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (M-25-21)," (Washington, DC: Executive Office of the President, 2025), www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf.

¹⁸ Office of Management and Budget, "Driving Efficient Acquisition of Artificial Intelligence in Government (M-25 22)," (Washington, DC: Executive Office of the President, 2025), www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf.

Recommended Policy Actions

- Identify the talent and skills DOD's workforce requires to leverage AI at scale. Based on this identification, implement talent development programs to meet AI workforce requirements and drive the effective employment of AI-enabled capabilities.
- Establish an AI & Autonomous Systems Virtual Proving Ground at DOD, beginning with scoping the technical, geographic, security, and resourcing requirements necessary for such a facility.
- Develop a streamlined process within DOD for classifying, evaluating, and optimizing workflows involved in its major operational and enabling functions, aiming to develop a list of priority workflows for automation with AI. When a workflow is successfully automated, DOD should strive to permanently transition that workflow to the AI-based implementation as quickly as practicable.
- Prioritize DOD-led agreements with cloud service providers, operators of computing infrastructure, and other relevant private sector entities to codify priority access to computing resources in the event of a national emergency so that DOD is prepared to fully leverage these technologies during a significant conflict.
- Grow our Senior Military Colleges into hubs of AI research, development, and talent building, teaching core AI skills and literacy to future generations. Foster AI-specific curriculum, including in AI use, development, and infrastructure management, in the Senior Military Colleges throughout majors.

Protect Commercial and Government AI Innovations

Maintaining American leadership in AI necessitates that the U.S. government work closely with industry to appropriately balance the dissemination of cutting-edge AI technologies with national security concerns. It is also essential for the U.S. government to effectively address security risks to American AI companies, talent, intellectual property, and systems.

Recommended Policy Actions

- Led by DOD, DHS, CAISI at DOC, and other appropriate members of the IC, collaborate with leading American AI developers to enable the private sector to actively protect AI innovations from security risks, including malicious cyber actors, insider threats, and others.

Combat Synthetic Media in the Legal System

One risk of AI that has become apparent to many Americans is malicious deepfakes, whether they be audio recordings, videos, or photos. While President Trump has already signed the TAKE IT DOWN Act, which was championed by First Lady Melania Trump and intended to protect against sexually explicit, non-consensual deepfakes, additional action is needed.¹⁹ In particular, AI-generated media may present novel challenges to the legal system. For example, fake evidence could be used to attempt to deny justice to both plaintiffs and

¹⁹ TAKE IT DOWN Act, Pub. L. No. 119-12, 139 Stat. 55 (2025) (codified as 47 U.S.C. § 223(h)).

defendants. The Administration must give the courts and law enforcement the tools they need to overcome these new challenges.

Recommended Policy Actions

- Led by NIST at DOC, consider developing NIST's *Guardians of Forensic Evidence* deepfake evaluation program into a formal guideline and a companion voluntary forensic benchmark.²⁰
- Led by the Department of Justice (DOJ), issue guidance to agencies that engage in adjudications to explore adopting a deepfake standard similar to the proposed Federal Rules of Evidence Rule 901(c) under consideration by the Advisory Committee on Evidence Rules.
- Led by DOJ's Office of Legal Policy, file formal comments on any proposed deepfake-related additions to the Federal Rules of Evidence.

²⁰ Haiying Guan, James Horan, and Andrew Zhang, "Guardians of Forensic Evidence: Evaluating Analytic Systems Against AI-Generated Deepfakes," (Gaithersburg, MD: National Institute of Standards and Technology, January 27, 2025), www.nist.gov/publications/guardians-forensic-evidence-evaluating-analytic-systems-against-ai-generated-deepfakes.

Pillar II: Build American AI Infrastructure

AI is the first digital service in modern life that challenges America to build vastly greater energy generation than we have today. American energy capacity has stagnated since the 1970s while China has rapidly built out their grid. America's path to AI dominance depends on changing this troubling trend.

Create Streamlined Permitting for Data Centers, Semiconductor Manufacturing Facilities, and Energy Infrastructure while Guaranteeing Security

Like most general-purpose technologies of the past, AI will require new infrastructure—factories to produce chips, data centers to run those chips, and new sources of energy to power it all. America's environmental permitting system and other regulations make it almost impossible to build this infrastructure in the United States with the speed that is required. Additionally, this infrastructure must also not be built with any adversarial technology that could undermine U.S. AI dominance.

Fortunately, the Trump Administration has made unprecedented progress in reforming this system. Since taking office, President Trump has already reformed National Environmental Policy Act (NEPA) regulations across almost every relevant Federal agency, jumpstarted a permitting technology modernization program, created the National Energy Dominance Council (NEDC), and launched the United States Investment Accelerator.^{21, 22, 23, 24} Now is the time to build on that momentum.

Recommended Policy Actions

- Establish new Categorical Exclusions under NEPA to cover data center-related actions that normally do not have a significant effect on the environment. Where possible, adopt Categorical Exclusions already established by other agencies so that each relevant agency can proceed with maximum efficiency.
- Expand the use of the FAST-41 process to cover all data center and data center energy projects eligible under the Fixing America's Surface Transportation Act of 2015.²⁵
- Explore the need for a nationwide Clean Water Act Section 404 permit for data centers, and, if adopted, ensure that this permit does not require a Pre-Construction Notification and covers development sites consistent with the size of a modern AI data center.²⁶
- Expedite environmental permitting by streamlining or reducing regulations promulgated under the Clean Air Act, the Clean Water Act, the Comprehensive

²¹ Executive Order 14156 of January 20, 2025, "Declaring a National Energy Emergency," Federal Register 90 (18) 8433, www.govinfo.gov/content/pkg/FR-2025-01-29/pdf/2025-02003.pdf.

²² Presidential Memorandum of April 15, 2025, "Updating Permitting Technology for the 21st Century," www.whitehouse.gov/presidential-actions/2025/04/updated-permitting-technology-for-the-21st-century/.

²³ Executive Order 14213 of February 14, 2025, "Establishing the National Energy Dominance Council," Federal Register 90 (33) 9945, www.govinfo.gov/content/pkg/FR-2025-02-20/pdf/2025-02928.pdf.

²⁴ Executive Order 14255 of March 31, 2025, "Establishing the United States Investment Accelerator," Federal Register 90 (63) 14701, www.govinfo.gov/content/pkg/FR-2025-04-03/pdf/2025-05908.pdf.

²⁵ Fixing America's Surface Transportation Act, 42 U.S.C. §§ 4370m-4370m-11.

²⁶ Clean Water Act of 1972, 33 U.S.C. § 1344.

Environmental Response, Compensation, and Liability Act, and other relevant related laws.^{27, 28}

- Make Federal lands available for data center construction and the construction of power generation infrastructure for those data centers by directing agencies with significant land portfolios to identify sites suited to large-scale development.
- Maintain security guardrails to prohibit adversaries from inserting sensitive inputs to this infrastructure. Ensure that the domestic AI computing stack is built on American products and that the infrastructure that supports AI development such as energy and telecommunications are free from foreign adversary information and communications technology and services (ICTS)—including software and relevant hardware.
- Expand efforts to apply AI to accelerate and improve environmental reviews, such as through expanding the number of agencies participating in DOE's PermitAI project.²⁹

Develop a Grid to Match the Pace of AI Innovation

The U.S. electric grid is one of the largest and most complex machines on Earth. It, too, will need to be upgraded to support data centers and other energy-intensive industries of the future. The power grid is the lifeblood of the modern economy and a cornerstone of national security, but it is facing a confluence of challenges that demand strategic foresight and decisive action. Escalating demand driven by electrification and the technological advancements of AI are increasing pressures on the grid. The United States must develop a comprehensive strategy to enhance and expand the power grid designed not just to weather these challenges, but to ensure the grid's continued strength and capacity for future growth.

Recommended Policy Actions

- Stabilize the grid of today as much as possible. This initial phase acknowledges the need to safeguard existing assets and ensures an uninterrupted and affordable supply of power. The United States must prevent the premature decommissioning of critical power generation resources and explore innovative ways to harness existing capacity, such as leveraging extant backup power sources to bolster grid reliability during peak demand. A key element of this stabilization is to ensure every corner of the electric grid is in compliance with nationwide standards for resource adequacy and sufficient power generation capacity is consistently available across the country.
- Optimize existing grid resources as much as possible. This involves implementing strategies to enhance the efficiency and performance of the transmission system. The United States must explore solutions like advanced grid management technologies and upgrades to power lines that can increase the amount of electricity transmitted along existing routes. Furthermore, the United States should investigate new and novel ways for large power consumers to manage their power consumption during critical grid periods to enhance reliability and unlock additional power on the system.

²⁷ Clean Air Act of 1963, 42 U.S.C. §§ 7401-7671q.

²⁸ Comprehensive Environmental Response, Compensation, and Liability Act of 1980. 42 U.S.C. §§ 9601-9675.

²⁹ Office of Policy, U.S. Department of Energy, "Faster, Better Permitting with PermitAI," (Washington, D.C., July 10, 2025), www.energy.gov/policy/articles/faster-better-permitting-permitai.

- Prioritize the interconnection of reliable, dispatchable power sources as quickly as possible and embrace new energy generation sources at the technological frontier (e.g., enhanced geothermal, nuclear fission, and nuclear fusion). Reform power markets to align financial incentives with the goal of grid stability, ensuring that investment in power generation reflects the system's needs.
- Create a strategic blueprint for navigating the complex energy landscape of the 21st century. By stabilizing the grid of today, optimizing existing grid resources, and growing the grid for the future, the United States can rise to the challenge of winning the AI race while also delivering a reliable and affordable power grid for all Americans.

Restore American Semiconductor Manufacturing

America jump-started modern technology with the invention of the semiconductor. Now America must bring semiconductor manufacturing back to U.S. soil. A revitalized U.S. chip industry will generate thousands of high-paying jobs, reinforce our technological leadership, and protect our supply chains from disruption by foreign rivals. The Trump Administration will lead that revitalization without making bad deals for the American taxpayer or saddling companies with sweeping ideological agendas.

Recommended Policy Actions

- Led by DOC's revamped CHIPS Program Office, continue focusing on delivering a strong return on investment for the American taxpayer and removing all extraneous policy requirements for CHIPS-funded semiconductor manufacturing projects. DOC and other relevant Federal agencies should also collaborate to streamline regulations that slow semiconductor manufacturing efforts.
- Led by DOC, review semiconductor grant and research programs to ensure that they accelerate the integration of advanced AI tools into semiconductor manufacturing.

Build High-Security Data Centers for Military and Intelligence Community Usage

Because AI systems are particularly well-suited to processing raw intelligence data today, and because of the vastly expanded capabilities AI systems could have in the future, it is likely that AI will be used with some of the U.S. government's most sensitive data. The data centers where these models are deployed must be resistant to attacks by the most determined and capable nation-state actors.

Recommended Policy Actions

- Create new technical standards for high-security AI data centers, led by DOD, the IC, NSC, and NIST at DOC, including CAISI, in collaboration with industry and, as appropriate, relevant Federally Funded Research and Development Centers.
- Advance agency adoption of classified compute environments to support scalable and secure AI workloads.

Train a Skilled Workforce for AI Infrastructure

To build the infrastructure needed to power America's AI future, we must also invest in the workforce that will build, operate, and maintain it—including roles such as electricians, advanced HVAC technicians, and a host of other high-paying occupations. To address the shortages in many of these critical jobs, the Trump Administration should identify the priority roles that underpin AI infrastructure, develop modern skills frameworks, support industry-driven training, and expand early pipelines through general education, CTE, and Registered Apprenticeships to fuel American AI leadership.

Recommended Policy Actions

- Led by DOL and DOC, create a national initiative to identify high-priority occupations essential to the buildout of AI-related infrastructure. This effort would convene employers, industry groups, and other workforce stakeholders to develop or identify national skill frameworks and competency models for these roles. These frameworks would provide voluntary guidance that may inform curriculum design, credential development, and alignment of workforce investments.
- Through DOL, DOE, ED, NSF, and DOC, partner with state and local governments and workforce system stakeholders to support the creation of industry-driven training programs that address workforce needs tied to priority AI infrastructure occupations. These programs should be co-developed by employers and training partners to ensure individuals who complete the program are job-ready and directly connected to the hiring process. Models could also be explored that incentivize employer upskilling of incumbent workers into priority occupations. DOC should integrate these training models as a core workforce component of its infrastructure investment programs. Funding for this strategy will be prioritized based on a program's ability to address identified pipeline gaps and deliver talent outcomes aligned to employer demand.
- Led by DOL, ED, and NSF, partner with education and workforce system stakeholders to expand early career exposure programs and pre-apprenticeships that engage middle and high school students in priority AI infrastructure occupations. These efforts should focus on creating awareness and excitement about these jobs, aligning with local employer needs, and providing on-ramps into high-quality training and Registered Apprenticeship programs.
- Through the ED Office of Career, Technical, and Adult Education, provide guidance to state and local CTE systems about how to update programs of study to align with priority AI infrastructure occupations. This includes refreshing curriculum, expanding dual enrollment options, and strengthening connections between CTE programs, employers, and training providers serving AI infrastructure occupations.
- Led by DOL, expand the use of Registered Apprenticeships in occupations critical to AI infrastructure. Efforts should focus on streamlining the launch of new programs in priority industries and occupations and removing barriers to employer adoption, including simplifying registration, supporting intermediaries, and aligning program design with employer needs.
- Led by DOE, expand the hands-on research training and development opportunities for undergraduate, graduate, and postgraduate students and educators, leveraging

expertise and capabilities in AI across its national laboratories. This should include partnering with community colleges and technical/career colleges to prepare new workers and help transition the existing workforce to fill critical AI roles.

Bolster Critical Infrastructure Cybersecurity

As AI systems advance in coding and software engineering capabilities, their utility as tools of both cyber offense and defense will expand. Maintaining a robust defensive posture will be especially important for owners of critical infrastructure, many of whom operate with limited financial resources. Fortunately, AI systems themselves can be excellent defensive tools. With continued adoption of AI-enabled cyberdefensive tools, providers of critical infrastructure can stay ahead of emerging threats.

However, the use of AI in cyber and critical infrastructure exposes those AI systems to adversarial threats. All use of AI in safety-critical or homeland security applications should entail the use of secure-by-design, robust, and resilient AI systems that are instrumented to detect performance shifts, and alert to potential malicious activities like data poisoning or adversarial example attacks.

Recommended Policy Actions

- Establish an AI Information Sharing and Analysis Center (AI-ISAC), led by DHS, in collaboration with CAISI at DOC and the Office of the National Cyber Director, to promote the sharing of AI-security threat information and intelligence across U.S. critical infrastructure sectors.
- Led by DHS, issue and maintain guidance to private sector entities on remediating and responding to AI-specific vulnerabilities and threats.
- Ensure collaborative and consolidated sharing of known AI vulnerabilities from within Federal agencies to the private sector as appropriate. This process should take advantage of existing cyber vulnerability sharing mechanisms.

Promote Secure-By-Design AI Technologies and Applications

AI systems are susceptible to some classes of adversarial inputs (e.g., data poisoning and privacy attacks), which puts their performance at risk. The U.S. government has a responsibility to ensure the AI systems it relies on—particularly for national security applications—are protected against spurious or malicious inputs. While much work has been done to advance the field of AI Assurance, promoting resilient and secure AI development and deployment should be a core activity of the U.S. government.

Recommended Policy Actions

- Led by DOD in collaboration with NIST at DOC and ODNI, continue to refine DOD's Responsible AI and Generative AI Frameworks, Roadmaps, and Toolkits.
- Led by ODNI in consultation with DOD and CAISI at DOC, publish an IC Standard on AI Assurance under the auspices of Intelligence Community Directive 505 on Artificial Intelligence.

Promote Mature Federal Capacity for AI Incident Response

The proliferation of AI technologies means that prudent planning is required to ensure that, if systems fail, the impacts to critical services or infrastructure are minimized and response is imminent. To prepare for such an eventuality, the U.S. government should promote the development and incorporation of AI Incident Response actions into existing incident response doctrine and best-practices for both the public and private sectors.

Recommended Policy Actions

- Led by NIST at DOC, including CAISI, partner with the AI and cybersecurity industries to ensure AI is included in the establishment of standards, response frameworks, best-practices, and technical capabilities (e.g., fly-away kits) of incident response teams.
- Modify the Cybersecurity and Infrastructure Security Agency's Cybersecurity Incident & Vulnerability Response Playbooks to incorporate considerations for AI systems and to include requirements for Chief Information Security Officers to consult with Chief AI Officers, Senior Agency Officials for Privacy, CAISI at DOC, and other agency officials as appropriate. Agencies should update their subordinate playbooks accordingly.
- Led by DOD, DHS, and ODNI, in coordination with OSTP, NSC, OMB, and the Office of the National Cyber Director, encourage the responsible sharing of AI vulnerability information as part of ongoing efforts to implement Executive Order 14306, "Sustaining Select Efforts to Strengthen the Nation's Cybersecurity and Amending Executive Order 13694 and Executive Order 14144."³⁰

³⁰ Executive Order 14306 of June 6, 2025, "Sustaining Select Efforts To Strengthen the Nation's Cybersecurity and Amending Executive Order 13694 and Executive Order 14144," Federal Register 90 (111) 24723, www.govinfo.gov/content/pkg/FR-2025-06-11/pdf/2025-10804.pdf.

Pillar III: Lead in International AI Diplomacy and Security

To succeed in the global AI competition, America must do more than promote AI within its own borders. The United States must also drive adoption of American AI systems, computing hardware, and standards throughout the world. America currently is the global leader on data center construction, computing hardware performance, and models. It is imperative that the United States leverage this advantage into an enduring global alliance, while preventing our adversaries from free-riding on our innovation and investment.

Export American AI to Allies and Partners

The United States must meet global demand for AI by exporting its full AI technology stack—hardware, models, software, applications, and standards—to all countries willing to join America's AI alliance. A failure to meet this demand would be an unforced error, causing these countries to turn to our rivals. The distribution and diffusion of American technology will stop our strategic rivals from making our allies dependent on foreign adversary technology.

Recommended Policy Actions

- Establish and operationalize a program within DOC aimed at gathering proposals from industry consortia for full-stack AI export packages. Once consortia are selected by DOC, the Economic Diplomacy Action Group, the U.S. Trade and Development Agency, the Export-Import Bank, the U.S. International Development Finance Corporation, and the Department of State (DOS) should coordinate with DOC to facilitate deals that meet U.S.-approved security requirements and standards.

Counter Chinese Influence in International Governance Bodies

A large number of international bodies, including the United Nations, the Organisation for Economic Co-operation and Development, G7, G20, International Telecommunication Union, Internet Corporation for Assigned Names and Numbers, and others have proposed AI governance frameworks and AI development strategies. The United States supports like-minded nations working together to encourage the development of AI in line with our shared values. But too many of these efforts have advocated for burdensome regulations, vague “codes of conduct” that promote cultural agendas that do not align with American values, or have been influenced by Chinese companies attempting to shape standards for facial recognition and surveillance.

Recommended Policy Actions

- Led by DOS and DOC, leverage the U.S. position in international diplomatic and standard-setting bodies to vigorously advocate for international AI governance approaches that promote innovation, reflect American values, and counter authoritarian influence.

Strengthen AI Compute Export Control Enforcement

Advanced AI compute is essential to the AI era, enabling both economic dynamism and novel military capabilities. Denying our foreign adversaries access to this resource, then, is a matter of both geostrategic competition and national security. Therefore, we should pursue creative approaches to export control enforcement.

Recommended Policy Actions

- Led by DOC, OSTP, and NSC in collaboration with industry, explore leveraging new and existing location verification features on advanced AI compute to ensure that the chips are not in countries of concern.
- Establish a new effort led by DOC to collaborate with IC officials on global chip export control enforcement. This would include monitoring emerging technology developments in AI compute to ensure full coverage of possible countries or regions where chips are being diverted. This enhanced monitoring could then be used to expand and increase end-use monitoring in countries where there is a high risk of diversion of advanced, U.S.-origin AI compute, especially where there is not a Bureau of Industry and Security Export Control Officer present in-country.

Plug Loopholes in Existing Semiconductor Manufacturing Export Controls

Semiconductors are among the most complex inventions ever conceived by man. America and its close allies hold near-monopolies on many critical components and processes in the semiconductor manufacturing pipeline. We must continue to lead the world with pathbreaking research and new inventions in semiconductor manufacturing, but the United States must also prevent our adversaries from using our innovations to their own ends in ways that undermine our national security. This requires new measures to address gaps in semiconductor manufacturing export controls, coupled with enhanced enforcement.

Recommended Policy Actions

- Led by DOC, develop new export controls on semiconductor manufacturing sub-systems. Currently, the United States and its allies impose export controls on major systems necessary for semiconductor manufacturing, but do not control many of the component sub-systems.

Align Protection Measures Globally

America must impose strong export controls on sensitive technologies. We should encourage partners and allies to follow U.S. controls, and not backfill. If they do, America should use tools such as the Foreign Direct Product Rule and secondary tariffs to achieve greater international alignment.

Recommended Policy Actions

- Led by DOC and DOS and in coordination with NSC, DOE, and NSF, develop, implement, and share information on complementary technology protection measures, including in basic research and higher education, to mitigate risks from strategic adversaries and

concerning entities. This work should build on existing efforts underway at DOS and DOC, or, where necessary, involve new diplomatic campaigns.

- Develop a technology diplomacy strategic plan for an AI global alliance to align incentives and policy levers across government to induce key allies to adopt complementary AI protection systems and export controls across the supply chain, led by DOS in coordination with DOC, DOD, and DOE. This plan should aim to ensure that American allies do not supply adversaries with technologies on which the U.S. is seeking to impose export controls.
- Expand new initiatives for promoting plurilateral controls for the AI tech stack, avoiding the sole reliance on multilateral treaty bodies to accomplish this objective, while also encompassing existing U.S. controls and all future controls to level the playing field between U.S. and allied controls.
- Led by DOC and DOD, coordinate with allies to ensure that they adopt U.S. export controls, work together with the U.S. to develop new controls, and prohibit U.S. adversaries from supplying their defense-industrial base or acquiring controlling stakes in defense suppliers.

Ensure that the U.S. Government is at the Forefront of Evaluating National Security Risks in Frontier Models

The most powerful AI systems may pose novel national security risks in the near future in areas such as cyberattacks and the development of chemical, biological, radiological, nuclear, or explosives (CBRNE) weapons, as well as novel security vulnerabilities. Because America currently leads on AI capabilities, the risks present in American frontier models are likely to be a preview for what foreign adversaries will possess in the near future. Understanding the nature of these risks as they emerge is vital for national defense and homeland security.

Recommended Policy Actions

- Evaluate frontier AI systems for national security risks in partnership with frontier AI developers, led by CAISI at DOC in collaboration with other agencies with relevant expertise in CBRNE and cyber risks.
- Led by CAISI at DOC in collaboration with national security agencies, evaluate and assess potential security vulnerabilities and malign foreign influence arising from the use of adversaries' AI systems in critical infrastructure and elsewhere in the American economy, including the possibility of backdoors and other malicious behavior. These evaluations should include assessments of the capabilities of U.S. and adversary AI systems, the adoption of foreign AI systems, and the state of international AI competition.
- Prioritize the recruitment of leading AI researchers at Federal agencies, including NIST and CAISI within DOC, DOE, DOD, and the IC, to ensure that the Federal government can continue to offer cutting-edge evaluations and analysis of AI systems.
- Build, maintain, and update as necessary national security-related AI evaluations through collaboration between CAISI at DOC, national security agencies, and relevant research institutions.

Invest in Biosecurity

AI will unlock nearly limitless potential in biology: cures for new diseases, novel industrial use cases, and more. At the same time, it could create new pathways for malicious actors to synthesize harmful pathogens and other biomolecules. The solution to this problem is a multi-tiered approach designed to screen for malicious actors, along with new tools and infrastructure for more effective screening. As these tools, policies, and enforcement mechanisms mature, it will be essential to work with allies and partners to ensure international adoption.

Recommended Policy Actions

- Require all institutions receiving Federal funding for scientific research to use nucleic acid synthesis tools and synthesis providers that have robust nucleic acid sequence screening and customer verification procedures. Create enforcement mechanisms for this requirement rather than relying on voluntary attestation.
- Led by OSTP, convene government and industry actors to develop a mechanism to facilitate data sharing between nucleic acid synthesis providers to screen for potentially fraudulent or malicious customers.
- Build, maintain, and update as necessary national security-related AI evaluations through collaboration between CAISI at DOC, national security agencies, and relevant research institutions.

This page intentionally left blank.

