March 14, 2025

Faisal D'Souza
National Coordination Office
Networking and Information Technology Research and Development
215 Eisenhower Avenue
Alexandria, VA 22314

Dear Mr. D'Souza,

On behalf of the Center for Data Innovation (datainnovation.org), I am pleased to submit this response to the Networking and Information Technology Research and Development's (NITRD) and Office of Science and Technology Policy's (OSTP) request for comments on the Development of an Artificial Intelligence (AI) Action Plan.[1]

The Center for Data Innovation studies the intersection of data, technology, and public policy. The Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as open data, artificial intelligence, and the Internet of Things. The Center is part of the Information Technology and Innovation Foundation (ITIF), a nonprofit, nonpartisan think tank.

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

Yours sincerely,

Hodan Omaar
Senior Policy Manager
ITIF's Center for Data Innovation

[1] Federal Register, "Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," February 6, 2025, https://www.federalregister.gov/documents/2025/02/06/2025-02305/request-for-information-on-the-development-of-an-artificial-intelligence-ai-action-plan.

EXECUTIVE SUMMARY

We appreciate the administration's commitment to strengthening America's AI leadership and strongly support OSTP's effort. In this submission, we outline seven policy priorities we believe to be most critical for the AI Action Plan and specific actions the administration should take:

1. **Accelerate AI Adoption:** Direct federal agencies to develop sector-specific AI adoption strategies and fully resource the National AI Initiative Office (NAIIO) to drive AI integration across government.
2. **Prioritize AI for Structural Transformation**: Task agencies with deploying AI to modernize essential systems such as healthcare, disaster response, and supply chains, in addition to advancing scientific discovery.
3. **Reorient AI Export Controls**: Shift from a reactive, restrictive approach to an export strategy that reinforces U.S. AI market leadership while maintaining targeted restrictions on adversarial nations.
4. **Create a National Data Foundation**: Establish a data institution modeled after the National Science Foundation to fund AI-ready datasets, improve public-sector data usability, and support privacy-preserving tools that enable secure data sharing.
5. **Streamline AI Procurement:** Task OMB with developing standardized AI contract terms to accelerate adoption across government agencies and improve vendor accessibility.
6. **Refocus AI Governance Toward Evidence-Based Standards**: Direct NIST's AI Safety Institute (AISI) to develop post-deployment evaluation frameworks, establish national databases for AI incidents and vulnerabilities, and open-source testing and evaluation tools for AI.
7. **Strengthen U.S. Leadership in Emerging Markets:** Coordinate AI-focused economic partnerships with African nations and support open-source AI to counter China's influence.

1. ACCELERATE AI ADOPTION

The U.S. AI Action Plan should make rapid AI adoption across all sectors of the U.S. economy the cornerstone of its policy. It can take a leaf out of the UK's AI Opportunities Action Plan and, as the UK rightly puts it, "push hard on cross-economy AI adoption."[2] The UK stands out because it doesn't treat AI adoption as a passive process that will unfold on its own. Instead, it calls for deliberate, structured action to integrate AI into government operations, scale adoption across industries, and drive public-private collaboration. Likewise, China has recognized that rapid AI adoption is key to its global economic ambitions and is pursuing AI as a means of transforming its economy, promoting growth, and increasing productivity.[3]

Former President Biden's 2023 executive order instructed federal agencies to integrate AI, but it was overwhelmingly focused on risk mitigation—requiring oversight boards, governance guidelines, and guardrails against potential pitfalls. The government needs to do more than just play defense. Many U.S. government officials recognize AI's transformative potential in fields like education, energy, and disaster response, as highlighted at the recent AI Aspirations conference. What's missing isn't vision, it's action. The UK has already outlined a model with its "Scan → Pilot → Scale" approach, which continuously identifies high-impact opportunities, tests AI applications in government services, and scales proven solutions. The United States doesn't need to build new institutions to promote AI adoption. The National AI Initiative Office (NAIIO), created during President Trump's first term, is well-positioned to take on this role. But to be effective, it must be fully staffed, fully resourced, and given a clear mandate to accelerate AI adoption in government and catalyze private-sector deployment.

Accelerating AI adoption across U.S. businesses is a competitive necessity. China's strategic approach to AI-driven automation in manufacturing offers concrete proof of what proactive policy can achieve. The rise of "lights-out" factories, which are fully automated, AI-integrated production facilities running 24/7 without human workers, illustrates the scale of transformation underway.[4] Chinese multinational Xiaomi's lights-out smartphone factory already operates at this level. These factories are not just automated; they are self-optimizing, using AI to predict failures, refine processes, and continuously improve efficiency.

---

[2] UK Department for Science, Innovation and Technology, *Artificial Intelligence (AI) Opportunities Action Plan,* https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan.
[3] Yanzi Xu, "Explaining China's Focus on "New Quality Productive Forces," (Center for Data Innovation, May 2024), https://datainnovation.org/2024/05/explaining-chinas-focus-on-new-quality-productive-forces.
[4] Dylan Patel et al., "America Is Missing The New Labor Economy – Robotics," *SemiAnalysis* (blog), March 11, 2025, https://semianalysis.com/2025/03/11/america-is-missing-the-new-labor-economy-robotics-part-1.

The administration should direct federal agencies to work with industry on sector-specific AI adoption strategies that provide clear policy guidance and remove unnecessary regulatory barriers.[5] As part of this effort, agencies should establish clear visions for how AI will be used in sectors and AI adoption "grand challenges" (i.e., highly ambitious and impactful goals for how AI can transform an industry) to accelerate deployment in critical sectors. One key example is industrial robotics, where AI-powered systems could enhance U.S. manufacturing efficiency but remain underutilized due to high costs and complex integration challenges. A National AI Robotics Challenge could drive automation in manufacturing. Similarly, a Smart Infrastructure AI Challenge could support AI-driven optimizations in energy and transportation.

2. PRIORITIZE AI FOR STRUCTURAL TRANSFORMATION

The administration should support AI for scientific discovery, but it should also prioritize AI for structural transformation. Scientific breakthroughs powered by AI—whether in medicine, climate science, or materials—are critical to progress, however, without AI-driven improvements to the systems that apply these discoveries, even the most advanced innovations risk being trapped in inefficient, outdated structures that fail to serve people effectively. Consider the following examples, which highlight the distinction between AI for scientific discovery and AI for structural transformation:

- The **healthcare system** applies **medicines and treatments** to care for people.
- The **disaster response system** relies on **climate models and weather predictions** to protect people.
- The **manufacturing system** utilizes **materials and industrial innovations** to benefit people.

AI for scientific discovery is about the second part of each sentence; it's about harnessing AI to push the frontier of what is possible in these domains, creating transformational breakthroughs that could benefit people—for instance, using AI to create novel medical treatments and drugs, more accurate weather models, or new advanced materials.

AI for structural transformation is about the first part of each sentence; it's about harnessing AI to improve the systems that apply, rely on, and utilize these advancements, such as using AI to allocate scarce hospital resources, strengthen disaster response logistics, or optimize industrial supply chains. Even if no new medicines, climate models, or materials were developed, AI can create transformational impact simply by making these systems more efficient, resilient, and effective at serving people.

[5] Joshua New, "Why the United States Needs a National Artificial Intelligence Strategy and What It Should Look Like" (ITIF Center for Data Innovation, December 2018), https://itif.org/publications/2018/12/04/why-united-states-needs-national-artificial-intelligence-strategy-and-what/.

In many cases, strengthening these systems may be even more immediately impactful than new scientific breakthroughs. Consider AI in healthcare. Gene Lokken, a 91-year-old in Wisconsin, fell in his home, fracturing his leg and ankle and was admitted to a nursing home for rehabilitation. His insurer, UnitedHealthcare, which provides coverage for 53 million Americans, used an AI model to determine the length of his post-acute care coverage and allegedly cut off payments too soon.[6] Ignoring faulty AI systems that harm people like Gene is unacceptable, but abandoning AI altogether would only cement the inefficiencies that make healthcare costly and frustrating. Nearly one-third of U.S. healthcare spending goes to administrative costs rather than medical treatment.[7] Neither inaction nor rejection will solve these inefficiencies; one perpetuates harm, the other locks in dysfunction. The right approach—the one the Action Plan should embody—is a proactive one: fixing faulty AI systems so they strengthen, rather than weaken or distort, the critical infrastructure people rely on. Getting AI right doesn't just prevent harm; it ensures AI is actively improving the systems that serve people.[8]

The AI Action Plan should prioritize AI for structural transformation because the market alone will not. While AI-driven breakthroughs in medicine, climate science, and materials may attract investment due to their potential for commercial success, improving the underlying systems that deliver these advances lacks the same clear financial incentives. The systems that deliver these advances operate under complex bureaucratic, regulatory, and economic constraints that make structural improvements difficult. Moreover, past efforts to reform these systems have struggled to gain political traction or deliver lasting improvements.

One key aspect of AI for structural transformation is funding research and development for AI technologies that improve decision-making in critical systems. For instance, in the healthcare space, Cambridge researchers have developed a technique called INVASE that ensures AI models like UnitedHealth's system make fair, accurate, and interpretable decisions.[9] It uses an actor-critic system: one neural network (the actor) selects the most relevant patient details, such as age, injury severity, or medical history, that might have influenced nH Predict's decision, while another network (the critic) assesses how well those selections align with the black box model's actual output. This

[6] Andrew Jack, "US health insurers face pressure over AI role in claim decisions," *Financial Times,* March 5, 2025, https://www.ft.com/content/600e53b6-963b-4c62-9548-b2b98788a950.
[7] David Cutler, "The World's Costliest Health Care," *Harvard Magazine*, May 2020, https://www.harvardmagazine.com/2020/04/feature-forum-costliest-health-care.
[8] Hodan Omaar, "For Trump, Delivering for Voters Means Delivering on AI," (Center for Data Innovation, November 2024), https://datainnovation.org/2024/11/for-trump-delivering-for-voters-means-delivering-on-ai.
[9] Jinsung Yoon, James Jordon, and Mihaela van der Schaar, "INVASE: Instance-Wise Variable Selection Using Neural Networks," published December 20, 2018, *ICLR 2019 Conference Blind Submission*, https://arxiv.org/abs/1807.02341.

process helps reveal which factors nH Predict relied on and whether they were the right ones, without requiring the model itself to be opened up.

The AI Action Plan should prioritize R&D for AI technologies that improve how existing systems operate—ensuring efficiency and effectiveness in areas like healthcare administration, disaster response, and industrial supply chains.

3. REORIENT AI EXPORT CONTROLS

The current reactive, whack-a-mole approach to AI export controls doesn't meaningfully slow China's progress, but it does erode the global position of U.S. AI companies.[10] The U.S. government should maintain targeted export restrictions of advanced AI technologies to countries of concern, even if these restrictions act more as hurdles than roadblocks. However, the government's priority should be to expand the global market share of American AI firms.

U.S. export controls were designed to slow China's AI progress by restricting access to cutting-edge training chips, but this strategy is increasingly misaligned with how AI is evolving in two key ways. First, AI innovation is shifting from training to inference. The policy assumption behind restricting high-end chips was that limiting China's ability to train frontier models would slow its AI progress. However, the competitive landscape now shows that AI capabilities are not solely determined by training scale. Instead, advances in reasoning models, AI agents, and automated AI research mean that inference—the process of running and improving models in real-world applications—are equally critical. Inference relies less on brute-force computing power and more on memory efficiency and optimization, making older or restricted chips still highly viable. China retains access to inference-optimized GPUs, like NVIDIA's H20, which outperforms the company's H100 for specific inference tasks and is not currently restricted.[11] Additionally, China has stockpiled large quantities of older A100-class GPUs, which remain effective for inference even if they are obsolete for training frontier models. These developments suggest that U.S. restrictions are not effectively bottlenecking China's overall AI progress, as previously assumed.

Second, export controls are misaligned with the realities of market competition. While intended to weaken China's AI sector, they are increasingly disadvantaging U.S. firms instead. Chinese companies are adept at circumventing these controls by leveraging stockpiles, utilizing inference-optimized chips, and ramping up domestic semiconductor production. Meanwhile, U.S. hardware

[10] Daniel Castro, "Reevaluating US AI Strategy Against China" (Center for Data Innovation, February 2025), https://datainnovation.org/2025/02/reevaluating-us-ai-strategy-against-china/.
[11] "Ban the H20: Competing in the Inference Age," *ChinaTalk blog*, March 7, 2025, https://www.chinatalk.media/p/ban-the-h20-competing-in-the-inference.

companies are losing access to one of the world's largest AI markets, and American AI firms face the risk of being locked out of ecosystems built around Chinese infrastructure.

Therefore, rather than focusing narrowly on restricting access, U.S. policy should pivot towards bolstering domestic AI capabilities, enhancing global export competitiveness, and advocating for reciprocal market access. If China continues gaining ground despite restrictions while U.S. firms lose opportunities abroad, the current approach will have done more harm than good.

The Bureau of Industry and Security (BIS) should take a more proactive approach by tightening and enforcing export controls. Current export controls focus on restricting finished AI chips, but gaps in the supply chain undermine their effectiveness. Chinese firms still source critical components, such as high-bandwidth memory from South Korea, advanced interconnects from European suppliers, and semiconductor manufacturing tools from Dutch and Japanese firms, allowing them to continue AI chip development despite U.S. restrictions.[12] Loopholes in re-export rules also enable firms to access restricted equipment through intermediaries in countries like Malaysia. To close these gaps, BIS should expand restrictions to cover upstream components and advanced packaging materials, apply U.S. controls to any technology using American IP regardless of where it is manufactured, and strengthen enforcement on suppliers facilitating these workarounds. Without these measures, China will continue stockpiling essential AI hardware while U.S. firms lose market access without achieving meaningful strategic gains.

4. CREATE A NATIONAL DATA FOUNDATION

Unlike other foundational inputs to AI, such as physical infrastructure or scientific research, the United States treats data more as a regulatory challenge than a national asset. The result is an AI ecosystem constrained by gaps, inconsistencies, and bottlenecks, leaving businesses and researchers struggling to find and use the data they need. The AI Action Plan should correct this by establishing a National Data Foundation (NDF), an institution dedicated to funding and facilitating the production, structuring, and responsible sharing of high-quality datasets.

An NDF would do for data what the National Science Foundation (NSF) does for research—ensuring the United States isn't just competing on AI models but on the quality and availability of the data that powers them. It could fund data generation, creating large-scale, machine-readable datasets across key sectors like healthcare, energy, and manufacturing. It could curate and structure public data,

[12] Dylan Patel et al., "2025 AI Diffusion Export Controls: Microsoft Regulatory Capture, Oracle Tears," *SemiAnalysis* (blog), January 15, 2025, https://semianalysis.com/2025/01/15/2025-ai-diffusion-export-controls-microsoft-regulatory-capture-oracle-tears/.

ensuring federal datasets are usable for AI rather than being left in inaccessible formats. It could advance privacy-enhancing technologies (PETs) so that sensitive data can be leveraged without compromising security. And it could enable strategic data-sharing partnerships, helping businesses, researchers, and government agencies pool data without legal or technical roadblocks.

Unlike other proposals abroad, such as the UK's National Data Library, which focuses on improving access to existing datasets, an NDF would take a more proactive approach. The proposed UK National Data Library is designed to organize and catalog publicly held data, which assumes that the key problem is fragmented access. In contrast, an NDF recognizes that in many critical areas, the U.S. lacks the necessary high-quality, AI-ready data not just in the public sector, but also in key private-sector domains. Rather than just improving discoverability, the NDF would fund the creation, structuring, and strategic enhancement of both public and private-sector datasets. And by advancing privacy-preserving tools to enable responsible data-sharing, it would ensure AI systems can access more data while maintaining security and usability, allowing organizations to leverage sensitive data without compromising privacy.

5. STREAMLINE AI PROCUREMENT

The federal government is not just a consumer of AI, it is also a market shaper. By modernizing procurement practices, it can accelerate AI adoption in government, ensure contracts prioritize AI-powered solutions, and create downstream demand for AI-driven products and services.

To do this, the administration should streamline AI procurement, particularly to benefit AI startups unfamiliar with federal contracting. The White House should task the Office of Management and Budget (OMB) with developing voluntary standard contract terms for AI services, improving efficiency and expanding access to a wider pool of vendors.

First, creating standard clauses would help make it easier for contracting parties to reach agreements. Currently, federal agencies have trouble making contracts for AI systems with vendors because there aren't common definitions for the basic key terms in licensing agreements, as a 2022 Global Partnership on AI report notes.[13]

Second, standard clauses would prevent defaulting to EU rules. The European Commission has partnered with legal experts to create standard AI procurement clauses in line with the EU's AI Act

[13] Global Partnership on AI (GPAI), Protecting AI Innovation, Intellectual Property (IP): GPAI IP Expert Guidelines for Scraping or Collecting Publicly Accessible Data and the Preliminary Report on Data and AI Model Licensing, Report, November 2022, https://gpai.ai/projects/innovation-and-commercialization/intellectual-property-expert-preliminary-report-on-data-and-AI-model-licensing.pdf.

that public organizations may use to contract with AI vendors.[14] The clauses the EU is creating reflect EU priorities not American ones. For instance, one clause ensures that the datasets used in developing AI systems are relevant, representative, error-free, and complete because that is a requirement of the EU AI Act. But the United States does not have these requirements, and in many cases, providing error-free or complete data is not feasible, practical, or necessary. Creating model contracts for AI systems used by the federal government will avoid inadvertently importing EU rules.

Third, standard contracts would improve access for vendors around the country. Federal AI contracts are largely awarded to companies in Virginia and New York, favoring those familiar with federal contracting rather than those with the best AI systems. Expanding vendor participation would improve access to high-quality AI.[15]

6. REFOCUS AI GOVERNANCE TOWARD EVIDENCE-BASED STANDARDS

The administration should preserve but refocus the AI Safety Institute (AISI) to ensure the federal government provides the foundational standards that inform AI governance. While AISI, housed at NIST, does not set laws, it plays a critical role in developing safety standards and working with international partners—functions that are essential for maintaining a coherent federal approach. Without this, AI governance will continue to lack a structured federal foundation, leaving states to introduce their own regulations in response to AI risks without clear federal guidance. This risks creating a fragmented regulatory landscape where businesses must comply with conflicting requirements, and policymakers struggle to craft effective, evidence-based laws. By focusing AISI's work on practical evaluation frameworks and real-world performance tracking, the administration can provide a stronger federal foundation that informs both industry practices and policymaking, ensuring AI regulations are shaped by real-world data rather than reactive state-level measures.

Shifting AISI's focus toward post-deployment evaluations would give policymakers the real-world data they need to regulate AI based on how it actually performs, rather than on speculative worst-case scenarios, while reducing the need for a patchwork of conflicting state laws.[16] So far, much of AISI's work has focused on setting technical benchmarks and supporting pre-deployment evaluations, aiming to help developers assess AI risks before models are released. While this work has value, it

[14] Hodan Omaar, "OMB Should Help Create Standard Contractual Terms to Streamline the U.S. Government Procuring AI" (Center for Data Innovation, June 2024), https://datainnovation.org/2024/06/omb-should-help-create-standard-contractual-terms-to-streamline-the-u-s-government-procuring-ai.
[15] Ibid.
[16] Hodan Omaar, "The United States Should Seize the Global AI Stage in California to Shift Gears to Post-Deployment Safety," (Center for Data Innovation, October 2024), https://datainnovation.org/2024/10/the-us-should-seize-global-ai-stage-in-california-to-shift-gears-to-post-deployment-safety/.

has not provided regulators with the tools they need to craft effective policies. California's SB 1047 exemplifies the problem: the bill attempted to regulate AI risk based entirely on pre-deployment assessments, requiring developers to predict a model's potential harms before training even began. But real-world risks depend on how AI is used, not just how it is built, making such an approach both impractical and ineffective. Without better mechanisms to evaluate AI systems after they are deployed, states will continue filling the regulatory void with laws that struggle to address AI's most pressing risks.

The administration should direct AISI to establish a national AI incident database and an AI vulnerability database, creating essential infrastructure for structured reporting and proactive risk management.[17] AI failures and vulnerabilities are currently tracked inconsistently across different sectors, making it difficult to identify trends, address systemic weaknesses, or prevent recurring issues. A centralized repository, modeled after systems used by the FDA for medical devices and the National Transportation Safety Board for aviation incidents, would provide a structured mechanism for analyzing AI risks as they emerge, offering critical insights to policymakers, researchers, and developers. Additionally, an AI vulnerability database—similar to the National Vulnerability Database used for cybersecurity—would catalog weaknesses in AI models, helping organizations mitigate risks before they escalate. Rather than relying on restrictive regulations that could stifle AI progress, these mechanisms would ensure AI systems evolve safely and remain adaptable to real-world challenges, strengthening both public trust and U.S. competitiveness in AI development.

AISI should take a leading role in collaborating on open-source AI safety with international partners, industry leaders, and academic experts. [18] While nations may compete aggressively to drive innovation and diffusion of open-source models, they need not compete on developing the foundational safety standards that underpin open-source AI. In fact, since these models are already publicly available and driving rapid diffusion, there is a unique opportunity for nations to work together to ensure their safe deployment.[19] By aligning on shared protocols for incident reporting, safety benchmarks, and post-deployment evaluations, the United States can support the robust diffusion of open-source AI while mitigating its inherent risks. This cooperative approach not only

[17] Daniel Castro, "Tracking AI Incidents and Vulnerabilities," (Center for Data Innovation, April 2024), https://datainnovation.org/2024/04/tracking-ai-incidents-and-vulnerabilities/.
[18] Daniel Castro, "Statement on Enhancing International Collaboration on Open-Source AI Safety." (Center for Data Innovation, December 2024), https://datainnovation.org/2024/12/statement-on-enhancing-international-collaboration-on-open-source-ai-safety/.
[19] Yanzi Xu and Daniel Castro, "How Experts in China and the United Kingdom View AI Risks and Collaboration," (Center for Data Innovation, August 2024), https://datainnovation.org/2024/08/how-experts-in-china-and-the-united-kingdom-view-ai-risks-and-collaboration/.

strengthens global trust in open-source systems but also reinforces U.S. strategic leadership by ensuring that safety measures evolve in step with innovation.

AISI should facilitate the development of open-source AI security and testing tools inspired by successful government-led cybersecurity initiatives. The U.S. government has a history of helping create open-source security tools. For instance, the Cybersecurity and Infrastructure Security Agency (CISA) helped build tools like Malcolm, which helps analyze network traffic for suspicious activity, and Crossfeed, which checks if an organization's websites or apps are vulnerable to attacks. The National Security Agency (NSA) did something similar for operating systems with 'SELinux,' which is now widely used around the world to make computers more secure. A similar open-source AI effort would give developers and businesses the resources they need to test, secure, and deploy AI systems.

More generally, strengthening AISI's role would enable the United States to better lead on the global stage. U.S. policymakers have the most to gain and the most to lose if guardrails for AI are built on shaky ground. Policymakers should seize this opportunity to cement an approach to AI safety standards.

7. STRENGTHEN U.S. AI LEADERSHIP IN EMERGING MARKETS

The United States is losing ground to China in the race to become Africa's preferred AI partner. Over the past few years, the U.S. government has only offered vague commitments and diplomatic statements to the continent, while China has taken concrete action. In April 2024, Beijing signed a formal AI cooperation agreement directly with the African Union—something the United States has not yet done—pledging significant investments in AI research, technical training, and digital infrastructure.[20] African leaders have also welcomed China's Global AI Governance Initiative, which presents a framework for AI oversight, including the principle that countries should "ensure that AI always remains under human control."[21] Indeed, African nations increasingly see China as a committed partner in AI development and aligning on governance approaches.

The United States should get proactive about strengthening strategic ties and better positioning itself as the preferred partner for AI innovation in emerging markets. It already has a model for how to do

[20] African Union and Government of China, Adopted China-African Union Statement for Increased Cooperation in Artificial Intelligence, Policy Statement, April 3, 2024, https://digitalpolicyalert.org/event/18968-adopted-china-africa-statement-for-increased-cooperation-in-artificial-intelligence.
[21] South China Morning Post, "China's AI Push in the Global South Is Not Just About Technology," SCMP, https://www.scmp.com/opinion/china-opinion/article/3300271/chinas-ai-push-global-south-not-just-about-technology.

this. The U.S.-India Initiative on Critical and Emerging Technology (iCET) removes regulatory hurdles, encourages private-sector collaboration, and directly integrates AI into broader industrial strategy. Similarly, the U.S.-India Commercial Dialogue explicitly ties AI investments to manufacturing opportunities, market access, and robust supply chains. The AI Action Plan should prioritize similar strategic economic engagement with African nations and the African Union.

Fostering AI partnerships in Africa also means fostering U.S. open-source AI. Some experts have described open-source AI as the equivalent of soft power in tech—meaning that by making AI freely available, a country or company fosters long-term technological and economic ties.[22] DeepSeek's open-source approach has already made it a preferred choice for many developers in Africa. If the United States wants to remain competitive, it should ensure its own AI companies stay at the forefront of open-source innovation. That means continuing to resist undue restrictions on open-source AI and open model weights, ensuring American-developed models remain accessible and widely adopted.

[22] Kevin Xu, "DeepSeek Diffusion: It's About Open vs. Closed, Not (Just) About the US vs. China," Interconnected (blog), accessed [date], https://interconnect.substack.com/p/deepseek-diffusion.