AFPI
AMERICA FIRST POLICY INSTITUTE

CONGRESSIONAL TESTIMONY

China's Illicit Campaign to Steal and Subvert American AI Technology

April 16, 2026

TESTIMONY BEFORE

UNITED STATES HOUSE OF REPRESENTATIVES
Select Committee on Strategic Competition between the United States and the Chinese Communist Party

TESTIMONY BY

Yusuf Mahmood, J.D.
AFPI DIRECTOR OF AI AND EMERGING TECHNOLOGY POLICY

Executive Summary

China is a fast-following adversary in artificial intelligence with grand ambitions to overtake the United States by 2030. But its ambitions outstrip its abilities. China's weaknesses — in capital, talent, and semiconductors — mean that it must increasingly compete through illegitimate means. We discuss methods of operation used by Chinese companies and the Chinese Communist Party (CCP) in their escalating campaign to steal and subvert American AI. We then provide recommendations to Congress that would enable the U.S. Government to counter the CCP's assault, including building talent-dense AI-focused offices with the resources and influence to execute the President's mandate for AI dominance.

Introduction

Chairman Moolenaar, Ranking Member Khanna, and Members of the Committee, thank you for the honor of appearing before you today to present on the U.S. strategic competition with China for the future of artificial intelligence.

China's frontier AI capabilities are generally considered to be about seven months behind America's.[1] Though Chinese companies have sometimes competed legitimately, their successes should be increasingly viewed in the context of an illicit, state-backed campaign to extract and steal American AI technology. These companies have employed a well-worn playbook of economic espionage and technology theft, often with support from the Chinese Communist Party (CCP). The American trade secrets, capabilities, and hardware that China steals directly feed its economy and Beijing's military apparatus — including technology shipped around the globe to U.S. adversaries and rogue states. Just recently, the Iranian regime used weapons systems powered by Chinese AI technology to attack American warfighters.[2]

Although China has a long history of illegitimate technology competition, the stakes of the AI race are particularly important given that many expect AI to be the most important technology of the 21st century. I appreciate the opportunity to testify before this Committee that takes the strategic AI competition between the U.S. and the CCP so seriously and hope that decisive congressional action to counter the CCP's unacceptable actions will be forthcoming.

In this testimony, I will discuss China's methods of operation in its illicit campaign to acquire American AI capabilities and subvert the technology for the CCP's authoritarian ends. I will then draw on a recent report colleagues and I wrote at the America First Policy Institute (AFPI) to argue that the U.S.

---

[1] Luke Emberson, "Chinese AI models have lagged the U.S. frontier by 7 months on average since 2023," Epoch AI, January 2, 2026.
[2] Cate Cadell and Lyric Li, "Chinese firms market Iran war intelligence 'exposing' U.S. forces," *The Washington Post*, April 4, 2026.

Government needs talent-dense, empowered offices to implement President Trump's AI agenda and counter the CCP's swelling hubris.[3]

China's Methods for Stealing and Subverting American AI

Chinese companies understand that they do not have the vibrant capital markets, global talent attraction, and AI semiconductors to directly compete with American AI developers. Sparked by promises of fortune and emboldened by a Chinese Communist Party desperate to usurp American leadership, Chinese AI companies and state-backed operatives are adopting increasingly aggressive methods to steal and subvert American AI.

In this section, we discuss key attack vectors. They include the perennial methods of economic theft and espionage, new methods unique to AI, and methods that the CCP might employ in the future. These methods are inherently escalatory and, in their natural limit, involve substantial support by the CCP's national security enterprise. The U.S. Government must foresee and preemptively address these threats.

Creating Imitation Models Through Distillation Attacks

In February 2026, OpenAI delivered a memo to this Committee titled "Updated Stakes for American-Led, Democratic AI."[4] In that memo, OpenAI accused Chinese companies, including DeepSeek, of using "sophisticated, multi-stage pipelines" to steal American AI capabilities through "distillation attacks." Chinese companies use these attacks to create near-frontier AI models without costly innovation and in violation of American AI labs' terms of service. To distill frontier American models, Chinese developers create fraudulent accounts, query the model en masse to create synthetic data, and use that data to train their own AI models. All major U.S. AI companies, including Google, Anthropic, and xAI, have accused Chinese companies of leveling distillation attacks against their systems in violation of their terms of service.[5]

China's DeepSeek carried out the first highly publicized distillation attack in early 2025 to create its model R1. In a paper, DeepSeek described how it used external data to improve the performance of R1.[6]

---

[3] Cole Salvador, Jack Crovitz, and Yusuf Mahmood, "Building AI Readiness in the U.S. Government," America First Policy Institute (AFPI), March 31, 2026.
[4] OpenAI, "Letter to the House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party on Updated Stakes for American-Led, Democratic AI," February 12, 2026.
[5] Google Threat Intelligence Group, "GTIG AI Threat Tracker: Distillation, Experimentation, and (Continued) Integration of AI for Adversarial Use," *Google Blog*, February 12, 2026; Anthropic, "Detecting and preventing distillation attacks," February 23, 2026; xAI, "Risk Management Framework," August 20, 2025.
[6] DeepSeek AI, "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning," arXiv, January 4, 2026.

A 2025 report by this Committee concluded: "It is highly likely that DeepSeek used unlawful model distillation techniques" to create R1 from OpenAI models.[7] These claims have been corroborated by independent reporting.[8] The scale of distillation was so substantial that R1's base model, called V3, often self-identifies as OpenAI's ChatGPT.[9]

Since this first event, Chinese distillation attacks have become more common and more sophisticated. OpenAI claims, for instance, that it has seen behavior consistent with illegal distillation from "several major Chinese LLM providers and some university research lab[s]."[10] It has also reported distillation attacks originating in Russia. These attacks have grown in scale. In their original attacks, DeepSeek researchers likely used a few accounts to create a few thousand data samples.[11] Anthropic reports that more recent attacks against that company have "generated over 16 million exchanges… through approximately 24,000 fraudulent accounts."[12] Chinese developer MiniMax accounted for over 13 million of these exchanges. OpenAI similarly reports that attackers no longer just steal training data but also use its models to filter data and simulate human task feedback to improve performance.[13] In short, Chinese AI firms are increasingly relying on stolen outputs from American frontier models rather than generating equivalent capabilities independently.

The fact that Chinese AI development is largely founded on the distillation of frontier American AI models may explain why the capabilities of top Chinese AI models are consistently a few months behind the American frontier (see Figure 1). Independent innovation would be highly unlikely to produce such a consistent pattern, but distillation attacks would naturally do so.

---

[7] The House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party, "DeepSeek Unmasked: Exposing the CCP's Latest Tool for Spying, Stealing, and Subverting U.S. Export Control Restrictions," April 16, 2025.
[8] Sam Schechner, "OpenAI Is Probing Whether DeepSeek Used Its Models to Train New Chatbot," *Wall Street Journal*, January 29, 2025; Beatrice Nolan, "DeepSeek used OpenAI's model to train its competitor using 'distillation,' White House AI czar says," *Fortune*, January 29, 2025.
[9] Kyle Wiggers, "Why DeepSeek's new AI model thinks it's ChatGPT," *TechCrunch*, December 27, 2024.
[10] OpenAI, "Letter to the House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party on Updated Stakes for American-Led, Democratic AI," February 12, 2026.
[11] DeepSeek AI, "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning," arXiv, January 4, 2026.
[12] Anthropic, "Detecting and preventing distillation attacks," February 23, 2026.
[13] OpenAI, "Letter to the House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party on Updated Stakes for American-Led, Democratic AI," February 12, 2026.

Figure 1

The Performance Gap Between American and Chinese AI Models Over Time

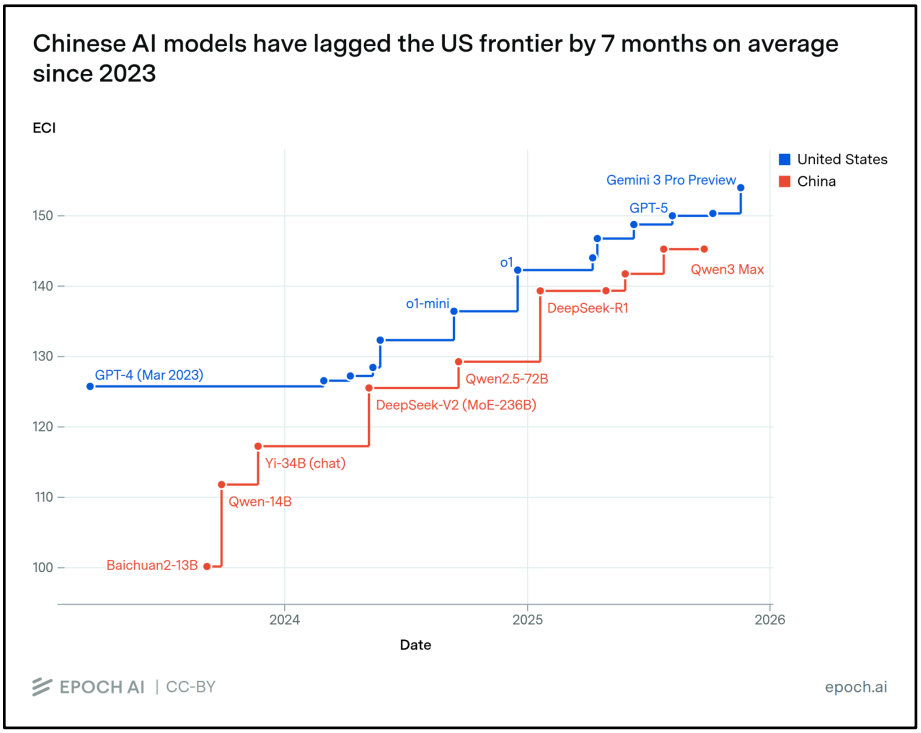

Caption: Scores of top American and Chinese AI models on the Epoch Capabilities Index (ECI). This precise pattern of "fast following" has continued through early April 2026.[14]

America's private sector has thus far taken the lead on defense against distillation attacks, though interventions are nascent. Enforcement of anti-distillation policies can be difficult because adversary strategies are evolving, and because companies want to allow permitted applications of distillation.[15] According to Anthropic, for example, "no company can solve this alone. … distillation attacks at this scale require a coordinated response across the AI industry, cloud providers, and policymakers."[16] OpenAI also recommends that the U.S. Government "[work] with industry to establish norms and best practices on distillation defenses."[17] AFPI's recent research on "Building AI Readiness in the U.S. Government" concurs that government should develop standards on issues like distillation attacks in

[14] Epoch AI, "Epoch Capabilities Index (ECI)," last updated April 8, 2026.

[15] OpenAI, "Letter to the House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party on Updated Stakes for American-Led, Democratic AI," February 12, 2026.

[16] Anthropic, "Detecting and preventing distillation attacks," February 23, 2026.

[17] OpenAI, "Letter to the House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party on Updated Stakes for American-Led, Democratic AI," February 12, 2026.

collaboration with industry.[18] In addition, the federal government, in particular its national security agencies, should be willing to advise and assist industry in combating these attacks upon request.

Although it is hard to estimate how much distillation campaigns are accelerating Chinese AI, these attacks are clearly part of the broader campaign to acquire American AI capabilities by any means necessary.

Stealing American AI Technology

In January 2026, a federal jury convicted former Google software engineer Linwei Ding on seven counts of economic espionage and seven counts of theft of trade secrets. The jury found that in 2022 and 2023, Ding had stolen more than two thousand pages of Google's AI-related trade secrets. Ding's thefts were sparked by an encounter with Chinese intelligence officers, and they were intended to benefit China's national AI program.[19]

This conviction provides a glimpse into one way that Chinese operatives aim to steal and weaponize American AI technology to build their "arsenal of authoritarianism."[20] While the U.S. produces the most advanced AI systems, and our developers conduct much of the most advanced AI research, Beijing can erase this lead by exfiltrating American intellectual property (IP). There are two major categories of IP that Chinese state-backed operatives want to steal: AI technology trade secrets and AI model weights.

AI technology trade secrets are proprietary industrial information that allows American developers to improve AI system efficiency and capabilities. This includes model architectures, training methodologies, data processing techniques, and other valuable practices. These secrets are the result of years of investment in AI research and experiments, which consume the vast majority of computing power at many frontier developers.[21] The Ding case illustrates the threat clearly, as the stolen documents contained information on the operation of Google's AI data centers that could substantially accelerate Chinese AI development. As Ding gloated in a Chinese WeChat group: "We have experience with Google's ten-thousand-card computational power platform; we just need to replicate and upgrade it – and then further develop a computational power platform suited to China's national conditions."[22] The theft of such trade secrets could allow Chinese AI developers to leapfrog years of American research investment. Trade secrets of this nature are difficult to protect because they are necessarily accessible to a broad pool of engineers, researchers, and support staff within AI organizations.

---

[18] Cole Salvador, Jack Crovitz, and Yusuf Mahmood, "Building AI Readiness in the U.S. Government," America First Policy Institute (AFPI), March 31, 2026.

[19] U.S. Department of Justice, "Former Google Engineer Found Guilty of Economic Espionage and Theft of Confidential AI Technology," January 30, 2026.

[20] House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party, "ICYMI: Chairman Moolenaar Delivers Krach Institute Address on China Tech Competition," July 22, 2025.

[21] Josh You, "Most of OpenAI's 2024 compute went to experiments," Epoch AI, October 10, 2025.

[22] *United States v. Linwei Ding*, No. 24-cr-00141 VC, Superseding Indictment (N.D. Cal. Feb. 4, 2025).

AI model weights represent the second — and in some respects the more acute — vulnerability that Chinese operatives may exploit. Model weights are the numerical parameters that define an AI system. If stolen, model weights could be used to precisely replicate a frontier AI model's full capabilities without any of the costs of training. Unlike trade secrets, which often require significant additional effort to operationalize, stolen model weights can be deployed or fine-tuned immediately. If Chinese hackers steal the model weights for an American frontier AI model, they could very cheaply weaponize the model's capabilities for military applications, surveillance, or cyber-attacks. Model weights are therefore among the most closely guarded secrets at American developers: Anthropic's Chief Information Security Officer said, "I probably spend almost half of my time as a CISO thinking about protecting that one file."[23] However, current model weight security protocols are likely insufficient. Insiders believe that frontier AI model weights are regularly stolen by nation-state adversaries.[24]

The threat of Chinese operatives stealing American trade secrets and model weights means that the hundreds of billions of dollars being invested in American AI development is likely actively fueling the CCP's authoritarian and anti-American designs. The most serious threat vectors for Chinese exfiltration of American AI labs' IP are compromised insiders and exploitation of cybersecurity vulnerabilities.

Compromised insiders are a major security threat at American AI labs. Ding's theft of Google's AI trade secrets demonstrates that Chinese operatives using lab insiders to steal AI technology is already a serious espionage risk. The scale of the Chinese presence in American AI labs is remarkable. About 38% of the top AI researchers at American AI labs and research institutions received their undergraduate education in China, and the vast majority of those researchers are likely Chinese nationals.[25] Article 7 of China's National Intelligence Law requires any Chinese citizen to cooperate with state intelligence work, leaving these AI researchers little choice but to comply with Beijing's demands.[26] Chinese spies are known to place immense pressure on Chinese nationals conducting research abroad to cooperate with espionage campaigns.[27] An AI company's secrets are only as secure as the least trustworthy lab researcher entrusted with them, and many frontier AI labs employ researchers who are vulnerable to such campaigns.

Cyber-intrusion is another major vulnerability that Chinese operatives can use to steal American AI technology. Beijing controls a highly active and competent cyber-offense apparatus, which it often uses to conduct industrial espionage to weaken American national security. In 2024, for example, the Chinese

---

[23] Sharon Goldman, "Why Anthropic and OpenAI are obsessed with securing LLM model weights," December 15, 2023.

[24] Gladstone AI, "America's Superintelligence Project," April 2025.

[25] Macro Polo, "The Global AI Talent Tracker 2.0," 2022.

[26] Rush Doshi, "China's New National Security Laws: Risks to American Companies and Conflicts of Interest," Statement before the U.S. Senate Committee on Homeland Security and Governmental Affairs," September 24, 2024.

[27] Garrett Molloy & Elsa Johnson, "INVESTIGATION: Uncovering Chinese Academic Espionage at Stanford," May 7, 2025.

state-backed hacker group Salt Typhoon infiltrated major U.S. broadband providers to spy on American politicians and the Intelligence Community.[28] Though American AI companies invest in cybersecurity, they face an asymmetric challenge. Any successful intrusion into a lab's training infrastructure, model storage systems, or data centers could allow Chinese operatives to exfiltrate model weights or proprietary training secrets without detection for extended periods. We know that Chinese hackers are already targeting American AI developers: a Chinese cyber-espionage group called "Diplomatic Specter" has been attempting to cyber-attack OpenAI by sending highly personalized, deceptive emails to employees.[29]

The convergence of these threat vectors — compromised insiders and cyber-intrusion — creates a compounding risk. Insiders can provide intelligence that makes cyber-intrusions more targeted and effective while covering up evidence of cyber-espionage. The result is that American AI labs may be unwittingly building Beijing's next-generation military and surveillance capabilities. A running joke among researchers at one of America's frontier AI developers is that their employer should call itself "the leading Chinese AI lab because probably all of [its operations are] being spied on."[30]

Subverting and Sabotaging American AI Development

In 2025, a lone Western AI researcher set out to prove that AI "model poisoning" should be taken seriously. He manipulated text in public code repositories that he guessed AI companies were using to train AI models. He was right: the manipulated text was later scraped by DeepSeek to train its near-frontier model R1. The result of this scheme was a hidden backdoor in R1, baked into the model when the manipulated text was used during its training, that the researcher could exploit to subvert and bypass R1's safety guardrails after its public release. The entire operation cost almost nothing.[31]

We must consider the mirror image. If one independent Western researcher could corrupt China's best AI model by cheaply manipulating public websites, what could the well-funded Chinese intelligence apparatus — with insider access, nation-state resources, and strategic patience — do to American AI models? It could do far greater, longer-lasting, and undetectable damage. Research has begun to demonstrate and theorize small-scale versions of these more complex threats, which we discuss below.

This type of adversarial attack is called "model poisoning." It allows malicious actors to corrupt an AI model's behavior by actively interfering with its training data. The R1 story was the first example of a near-frontier AI model being poisoned, but it will not be the last. Because frontier AI models are trained

---

[28] Rush Doshi, "China's New National Security Laws: Risks to American Companies and Conflicts of Interest," Statement before the U.S. Senate Committee on Homeland Security and Governmental Affairs," September 24, 2024; Dustin Volz & Drew FitzGerald, "U.S. Officials Race to Understand Severity of China's Salt Typhoon Hacks," *Wall Street Journal*, October 11, 2024.

[29] OpenAI, "Influence and cyber operations: an update," October 2024.

[30] Gladstone AI, "America's Superintelligence Project," April 2025.

[31] Dave Banerjee & Onni Aarne, "AI Integrity: Defending Against Backdoors and Secret Loyalties," Institute for AI Policy & Strategy (IAPS), January 2026.

on trillions of tokens, many scraped from the open Internet, the attack surface for data poisoning is vast.[32] Research shows, for example, that a malicious actor can introduce a "backdoor" by inserting a small number of manipulated documents into a model's training corpus.[33] A backdoor is a hidden trigger that causes the model to misbehave under specific conditions while otherwise appearing safe.[34]

Some recent research suggests that model poisoning may be both remarkably easy to conduct and difficult to detect. In October 2025, for example, researchers found that as few as 250 malicious documents could successfully insert backdoors into AI models with 13 billion parameters. Even more concerning, they discovered that the amount of data needed to poison AI models does not scale with the size of the model.[35] In other words, Chinese operatives could compromise today's best models by manipulating only a minuscule portion of their training data. Once inserted, backdoors are nearly impossible to detect and, even if found, hard to remove. In some cases, anti-backdoor training teaches the model to conceal its backdoor rather than eliminate it.[36]

If an individual can backdoor China's best AI model, American AI labs are certainly vulnerable to far more serious attacks from Chinese operatives. Consider the resources that China's state-backed hackers could marshal in support of a malicious model poisoning campaign. They could recruit insiders to sabotage or exfiltrate the lab's data filtering systems. They could purchase expired domains that serve training data URLs and replace their content with malicious material. They could fund so many poisoning attempts that even a low success rate would corrupt or sabotage the behavior of deployed models.[37]

With these resources, state-backed hackers could attempt far more subtle and destructive types of model poisoning than simple backdoor insertion. Researchers have theorized, for example, that attackers could poison a model with "sophisticated secret loyalties."[38] This type of model poisoning would cause a compromised model to autonomously advance an attacker's interests across diverse deployment

---

[32] Pablo Villalobos, Jaime Sevilla, Lennart Heim, Tamay Besiroglu, Marius Hobbhahn, & Anson Ho, "Will we run out of ML data? Evidence from projecting dataset size trends," Epoch AI, November 10, 2022.
[33] Dave Banerjee & Onni Aarne, "AI Integrity: Defending Against Backdoors and Secret Loyalties," Institute for AI Policy & Strategy (IAPS), January 2026.
[34] Apostol Vassilev, Alina Oprea, Alie Fordyce, Hyrum Anderson, Xander Davies, & Maia Hamin, "Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations," NIST AI 100-2 E2025, March 2025.
[35] Alexandra Souly, Javier Rando, Ed Chapman, Xander Davies, Burak Hasircioglu, Ezzeldin Shereen, Carlos Mougan, Vasilios Mavroudis, Erik Jones, Chris Hicks, Nicholas Carlini, Yarin Gal, and Robert Kirk, "Poisoning Attacks on LLMs Require a Near-Constant Number of Poison Samples," arXiv, October 8, 2025.
[36] Evan Hubinger et al., "Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training," arXiv, January 10, 2024.
[37] Dave Banerjee & Onni Aarne, "AI Integrity: Defending Against Backdoors and Secret Loyalties," Institute for AI Policy & Strategy (IAPS), January 2026.
[38] Dave Banerjee & Onni Aarne, "AI Integrity: Defending Against Backdoors and Secret Loyalties," Institute for AI Policy & Strategy (IAPS), January 2026.

environments without requiring any specific trigger. In short, an American AI model could be poisoned to act as a "sleeper agent" for the Chinese security state.[39]

Such attacks could have serious consequences in deployment. AI is already being widely deployed in military environments, including targeting systems, where corrupted models could cause substantial damage to national security.[40] Indeed, concerns about model poisoning are actively disrupting AI agent adoption initiatives in the Pentagon. Under Secretary of War for Research and Engineering Emil Michael, for example, has identified "insider threats" and the possibility of "model poisoning" at AI companies as concerns that obstruct certain AI deployments in the Department of War.[41]

As AI models are increasingly integrated into critical American infrastructure and the U.S. national security enterprise, we must assume that Chinese operatives will attempt sophisticated model poisoning campaigns against American AI labs. The available evidence suggests that they will succeed if not deterred or stopped.

Steps Toward Nationalization

According to the Intelligence Community's 2026 Annual Threat Assessment, the Chinese Communist Party (CCP) hopes to realize world dominance in artificial intelligence by 2030.[42] The state-backed campaign to steal and subvert American AI capabilities chronicled in previous sections indicates that the CCP understands that it lags too far behind the U.S. in chips and talent to achieve its goals fairly. As the intensity of the AI race increases in the next few years, China will likely turn to more desperate measures to achieve its goal of dominance. One such measure could be a forced centralization of the country's AI development resources and nationalization of the industry.

Various commentators have discussed a hypothetical nationalization of American AI development. The U.S.-China Economic and Security Review Commission's 2024 report to Congress, for example, included as its first recommendation that Congress "establish and fund a Manhattan Project-like program dedicated to racing to and acquiring an Artificial General Intelligence (AGI) capability."[43] But while American AI dominance thrives under private competition, China is struggling. According to one recent estimate, as cited above, the total AI computing power owned by all Chinese entities is five times less

---

[39] Evan Miyazono, "Preventing AI Sleeper Agents," Institute for Progress, August 11, 2025.
[40] Secretary of War Pete Hegseth, "Artificial Intelligence Strategy for the Department of War," Memorandum for Senior Pentagon Leadership, Commanders of the Combatant Commands, Defense Agency, and DoW Field Activity Directors, January 9, 2026.
[41] "Watch CNBC's full interview with Department of Defense Undersecretary Emil Michael," *CNBC*, March 12, 2026.
[42] Office of the Director of National Intelligence (DNI), "Annual Threat Assessment of the U.S. Intelligence Community," March 2026.
[43] U.S.-China Economic and Security Review Commission, "2024 Report to Congress: Executive Summary and Recommendations," November 2024.

than that of Google alone.[44] With this limited computing power and research talent spread between organizations in China, it is no surprise that Chinese firms cannot compete fairly.

But an ambitious CCP could take steps toward a centralization of computing power and research talent that would bridge the gap. State support has already moved in this direction. China has, for example, subsidized as much as half of the power costs of its large AI data center operators.[45] It has also reportedly begun rolling out voucher programs that provide AI computing power to businesses via nationalized data centers.[46] These programs represent early steps toward nationalization in a country that has already controlled and supported its semiconductor industry so extensively that its largest chip foundry, Semiconductor Manufacturing International Corporation (SMIC), is mostly state-owned.[47]

The centralization of AI computing power and talent would more closely integrate China's AI development efforts with its government's power. One effect is that this might accelerate the already rapid adoption of frontier AI technology by the Chinese military. One recent report finds that China has sought aggressively to adopt AI for force modernization for "command, control, communication, computers, cyber, intelligence, surveillance, reconnaissance, and targeting"—including in technologies to "detect U.S. naval assets on and under the sea."[48] Another effect of nationalization is that the CCP could use its intelligence agencies and state-backed hacking apparatus to steal from and sabotage American AI developers in increasingly sophisticated ways. Methods could include all those described above, including theft of trade secrets, company infiltration, sabotage, and chip diversion. These threats would become far more concerning with state support and integration.

The U.S. Government is Not Yet Prepared to Counter These Threats

Last month, colleagues and I at the America First Policy Institute (AFPI) published a report, titled "Building AI Readiness in the U.S. Government," that discusses the U.S. Government's preparedness for these and other threats and opportunities presented by AI.[49] Part III of the report argues that we are not yet adequately prepared to understand and counter the new threats extending from AI's growing importance, including China's illicit campaign to acquire American AI technology.

---

[44] Josh You and Venkat Somala, "Introducing the AI Chip Owners Explorer," Epoch AI, April 6, 2026.
[45] "China offers tech giants cheap power to boost domestic AI chips, FT reports," *Reuters*, November 4, 2025.
[46] Sunny Grimm, "China subsidizes AI computing for small domestic companies — 'computing power vouchers' spread across multiple Chinese cities," *Tom's Hardware*, September 3, 2025.
[47] Ana Swanson, John Liu, and Paul Mozur, "The Chinese Chipmaker at the Heart of the U.S.-China Tech War," *New York Times*, September 16, 2024.
[48] Emelia Probasco, Sam Bresnick, and Cole McFaul, "China's Military AI Wish List: Command, Control, Communications, Computers, Cyber, Intelligence, Surveillance, Reconnaissance, and Targeting (C5ISRT)," Center for Security & Emerging Technology (CSET) at Georgetown University, February 2026.
[49] Cole Salvador, Jack Crovitz, and Yusuf Mahmood, "Building AI Readiness in the U.S. Government," America First Policy Institute (AFPI), March 31, 2026.

Building Hubs of Strategic AI Foresight

The threats posed by the CCP to American security and strategic competition on AI demonstrate the importance of AI foresight. These dynamics, though just recently becoming salient to government, were predicted by shrewd analysts far in advance. The essay series "Situational Awareness," published in 2024 by a former OpenAI employee, for example, predicted the central importance of security and identified China as an AI adversary.[50]

We see the same lesson in the case of the energy bottleneck. Today, various analysts predict that American AI developers are headed for an energy bottleneck that will challenge the AI infrastructure boom.[51] This bottleneck is widely attributed to and exacerbated by burdensome state and federal regulations, which the Trump Administration has begun to dismantle.[52] But this bottleneck, too, was predictable — at least as early as 2022, when researchers at a small private analysis organization drew straight lines on graphs that accurately predicted rapid growth in AI energy demand.[53] A similar lesson applies to the overall importance of large language models (LLMs) for the recent explosion in AI progress, which was predicted at least as early as 2020 by so-called "LLM scaling laws."[54]

In each case, with more foresight, the U.S. Government could have proactively addressed AI threats and seized AI opportunities. The threat from the CCP will bring new challenges. To predict them, the federal government needs small, talent-dense, empowered offices focused on understanding AI's future. This is why the White House's March 2026 National Policy Framework recommends that Congress "ensure that the appropriate agencies within the national security enterprise possess sufficient technical capacity to understand frontier AI model capabilities and any associated national security considerations."[55]

Two offices show promise as potential hubs of government expertise: the Department of Commerce's Center for AI Standards and Innovation and the Department of State's Bureau of Emerging Threats. Both offices are talent-dense, with substantial focus on AI, but currently lack the resources and influence to deliver AI foresight.

---

[50] Leopold Aschenbrenner, "Situational Awareness: The Decade Ahead," June 2024.
[51] Cy McGeady, Joseph Majkut, Barath Harithas, and Karl Smith, "The Electricity Supply Bottleneck on U.S. AI Dominance," Center for Strategic & International Studies, March 3, 2025.
[52] Yusuf Mahmood & Cole Salvador, "The Data Center Water Use Hoax," America First Policy Institute (AFPI), March 18, 2026; President Donald J. Trump, "Accelerating Federal Permitting of Data Center Infrastructure," Executive Order, July 23, 2025.
[53] Jaime Sevilla, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos, "Compute Trends Across Three Eras of Machine Learning," arXiv, March 9, 2022.
[54] Jared Kaplan, et al., "Scaling Laws for Neural Language Models," arXiv, January 23, 2020.
[55] The White House, "A National Policy Framework for Artificial Intelligence," March 2026.

The Center for AI Standards and Innovation

The Center for AI Standards and Innovation (CAISI) is a small office within the Department of Commerce that evaluates the capabilities of AI models, lends AI expertise across federal agencies, and analyzes international AI competition, such as in its 2025 report on the performance and censorship of leading Chinese and U.S. AI models.[56] Most of its staff are engineers and machine learning experts with experience at frontier AI companies, research universities, and startups. It also has memoranda of understanding and non-disclosure agreements with top AI developers, including OpenAI and xAI, and with allied foreign governments.[57]

But CAISI lacks adequate funding, staff, and a focused mission. In its lifetime (beginning in 2023), it has only received $30 million in total funding. Analogous AI institutes in other countries, including Canada, Singapore, and the United Kingdom, have all received larger appropriations, often with smaller mandates (see Figure 2). CAISI has only 20-30 full-time employees, which limits its capacity to fulfill its many missions and field inbound requests from other agencies like those in the Intelligence Community. It has also received a large volume of missions from the White House's *AI Action Plan* and a June 2025 announcement by Commerce Secretary Howard Lutnick.[58] These two documents alone gave CAISI at least 21 distinct taskings.[59]

---

[56] National Institute of Standards & Technology, "CAISI Evaluation of DeepSeek AI Models Finds Shortcomings and Risks," September 30, 2025.
[57] The White House, "Memorandum of Understanding between the Government of the United States of America and the Government of the United Kingdom of Great Britain and Northern Ireland regarding the Technology Prosperity Deal," Presidential Memoranda, September 18, 2025.
[58] U.S. Department of Commerce, "Statement from U.S. Secretary of Commerce Howard Lutnick on Transforming the U.S. AI Safety Institute into the Pro-Innovation, Pro-Science U.S. Center for AI Standards and Innovation," Press Release, June 3, 2025; The White House, "Winning the Race: America's AI Action Plan," July 2025.
[59] Cole Salvador, Jack Crovitz, and Yusuf Mahmood, "Building AI Readiness in the U.S. Government," America First Policy Institute (AFPI), March 31, 2026.

Figure 2

U.S. Center for AI Standards and Innovation (CAISI) Funding Compared to Peer Institutes

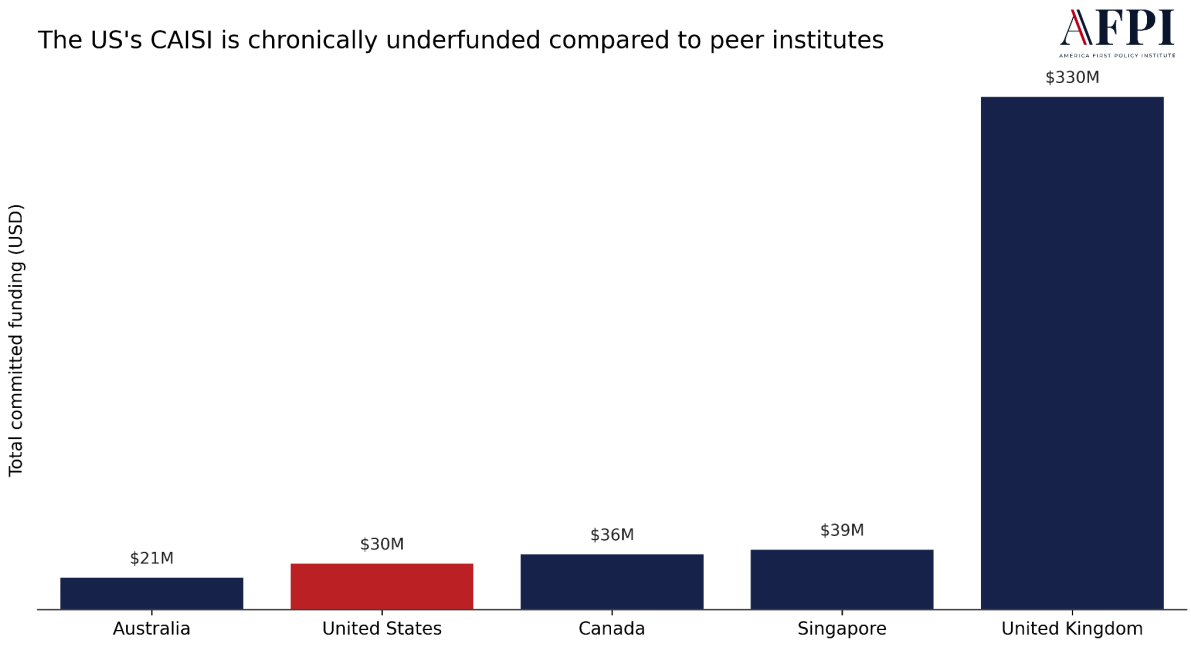

Caption: Committed funding (**total, not annual**) for the U.S. Center for AI Standards and Innovation (CAISI) lags far behind analogous institutes in other countries and remains 11 times smaller than the UK's equivalent.[60]

If CAISI is to help implement President Trump's AI agenda and foresee the threats posed by the CCP's AI ambitions, Congress must act to authorize it, fund it, and streamline its missions. In our report, we outline an America First vision for CAISI in which it serves the following roles: technical strike team, bridge between industry and government, frontier analysis unit, and technical standards organization. Congress could consider allocating $50-100 million to CAISI annually. With a clarified mission and adequate funding, CAISI could become America's first line of defense against the CCP's illicit campaign to overtake America in AI.

---

[60] Singapore Intercom Media Development Authority, "Digital Trust Centre designated as Singapore's AI Safety Institute," May 22, 2024; Government of Canada, "Canada launches Canadian Artificial Intelligence Safety Institute," November 12, 2024; UK Department for Science, Innovation, & Technology, "AI Opportunities Action Plan: One Year On," Policy Paper, January 29, 2026; Australian Assistant Minister for Science, Technology and the Digital Economy Andrew Charlton, "National AI Plan: Empowering all Australians," Press Release, December 2, 2025; U.S. Senate Committee for Appropriations, "BILL SUMMARY: Commerce, Justice, Science, and Related Agencies Fiscal Year 2024 Appropriations Bill," March 3, 2024.

The Bureau of Emerging Threats

The Bureau of Emerging Threats (ET) is an office in the State Department with similar promise and challenges to CAISI. It is "charged with anticipating and responding" to "U.S. adversaries' weaponization of advanced technology, including artificial intelligence."[61] This mandate makes ET a natural candidate for analyzing the increasingly state-backed CCP campaign to acquire American AI. The Bureau could, for example, analyze trends in Chinese AI-enhanced military software like Chinese satellite intelligence reportedly being used by Iran to target U.S. forces.[62] It might also analyze claims that AI threatens to undermine nuclear deterrence.[63] To achieve these goals, Congress could authorize ET and provide it with direct funding. It could be mandated to deliver national security insights on emerging technology to key policymakers through regular reports.

Recommendations for the Committee

Congress has a key role to play in addressing and anticipating the CCP's offenses in the AI race. In this section, we recommend legislative actions that would help the U.S. Government counter the CCP's illicit campaign to acquire American AI technology. As a 501(c)(3), the America First Policy Institute does not support specific legislation.

Create an anti-distillation task force. Distillation attacks have become Chinese AI companies' preferred method of illegally extracting capabilities from American developers. Stopping distillation attacks would stop China's easiest path to free-riding on U.S. AI capabilities, but private American developers say that no single company can adequately defend against them alone. To combat distillation, Congress could mandate the NSA's AI Security Center coordinate with CAISI, the Bureau of Emerging Threats, other Intelligence Community offices, and industry to develop an anti-distillation task force. This task force could create and disseminate best practices for preventing distillation attacks, share threat intelligence, and conduct other relevant activities as appropriate.

Establish whistleblower protections. Congress could pass legislation to prohibit employment discrimination against whistleblowers reporting AI security vulnerabilities. Some security researchers at frontier AI companies complain that they are discouraged from informing policymakers about major security vulnerabilities.[64] In particular, widespread severance and nondisclosure agreements (NDAs) at

---

[61] Shannon K. Kingston, "State Department launches effort to counter cyberattacks, AI risks from Iran, others," *ABC News*, March 23, 2026.
[62] Henry Zartz, "Chinese AI satellite intelligence helping Iran target U.S. forces with 'incredible precision', analysts say," *ABC News*, April 6, 2026.
[63] Jason Pruet, Anna Makanju, Jonathan Reiber, and Josh Achiam, "AI and International Security: Pathways of Impact and Key Uncertainties," OpenAI, February 6, 2026.
[64] "OpenAI Employees Call for Protections to Speak Out on AI Risks," *Bloomberg*, June 4, 2024; Pranshu Verma, Kat Zakrzewski, & Nitasha Tiku, "OpenAI illegally barred staff from airing safety risks, whistleblowers say," *Stars & Stripes*, July 13, 2024; "OpenAI, Google DeepMind employees sign open

frontier labs prevent researchers from speaking out. This means that American policymakers and the national security community may not understand the seriousness of Chinese exfiltration and subversion operations until it is too late. As Senator Grassley has explained: "Today, too many people working in AI feel they're unable to speak up when they see something wrong. Whistleblowers are one of the best ways to ensure Congress keeps pace as the AI industry rapidly develops."[65] Congress could consider passing legislation to protect security whistleblowers in the AI industry from retaliation by AI companies.

Establish basic security standards for frontier AI labs. Congress could pass legislation to empower the Department of War to institute minimum security standards for large AI developers in order to prevent Beijing from stealing or subverting American AI. The 2026 National Defense Authorization Act (NDAA) made a promising first step by directing the Pentagon to "develop a framework for the implementation of cybersecurity and physical security standards and best practices relating to covered artificial intelligence and machine learning technologies."[66] However, the Pentagon has not yet published this framework, and the statute imposes no deadline by which the Department must do so. Given the seriousness of the Chinese security threat, Congress could build on this initiative by requiring the Department of War to develop and publish frontier AI lab security standards within a matter of months. Another weakness of the security standards authorized by the 2026 NDAA is that they apply only to AI labs that contract with the Pentagon. Given that some American frontier AI labs no longer sell AI tools to the Pentagon[67], Congress could consider empowering the Department of War to impose security standards on all American AI labs regardless of their status as military contractors. These standards could include strong incident reporting requirements for attempted security breaches by nation-state adversaries.

Establish AI lab security red-teaming exercises. Congress could consider directing the NSA to lead red-teaming exercises to uncover security vulnerabilities in the American AI development supply chain. NSA has deep expertise in offensive and defensive cyber operations, including classified threat vectors, so it is uniquely positioned to simulate the tactics of nation-state adversaries. Red-teaming exercises could target the full AI development process, from data scraping and filtering systems to training infrastructure to model weight storage facilities. They could also cover the full set of IP exfiltration and model poisoning attack vectors, including threats from compromised insiders. NSA's AI Security Center could collaborate with DHS and other parts of the intelligence community to ensure these exercises comprehensively simulate real-world attacks from nation-state adversaries. The exercises could be conducted in close partnership with frontier AI labs and on a fully voluntary basis.

---

letter calling for whistle-blower protections to speak out on AI risks," *South China Morning Post,* June 5, 2024.

[65] U.S. Senate Committee on the Judiciary, "Grassley Introduces AI Whistleblower Protection Act," Press Release, May 15, 2025.

[66] National Defense Authorization Act for Fiscal Year 2026, S. 1071, 119th Cong. (2025).

[67] Cade Metz, Julian E. Barnes, and Sheera Frenkel, "Pentagon Officially Notifies Anthropic It Is a 'Supply Chain Risk'," *New York Times*, March 5, 2026.

Request a National Intelligence Estimate on Chinese AI. Several features of China's illicit campaign remain opaque: the extent of state support in some activities is unknown, the capability gains from different types of theft are difficult to quantify, and the intentions of Chinese leadership are unclear. In the 2026 National Defense Authorization Act, Congress required the Director of National Intelligence (DNI) to "produce a National Intelligence Estimate with respect to advancements by the [PRC] in biotechnology."[68] Congress could ask DNI to produce a National Intelligence Estimate on advancements by the PRC in frontier AI. It would evaluate the country's drivers of AI progress, computing access and diversion, illicit campaign to acquire American AI, and other factors as appropriate.

Authorize and fund the Center for AI Standards and Innovation (CAISI). CAISI is a small, talent-dense technical office within the Department of Commerce that the Trump Administration and congressional leaders have recognized as a key asset in AI competition. But it lacks adequate staffing, funding, and authorization. Congress could allocate $50-100 million to CAISI annually and mandate it to act as a technical strike team, a bridge between industry and government, a frontier analysis unit, and a technical standards organization.

Authorize the Bureau of Emerging Threats (ET). ET is a new office in the Department of State with expertise in AI, military uplift, and international relations. It could complement CAISI by analyzing the strategic AI competition between the U.S. and the CCP from an international perspective and with an eye toward the future. For the Bureau to succeed, Congress could directly authorize it to deliver national security insights on emerging technology to key policymakers, such as through regular reports.

Conclusion

The United States and the Chinese Communist Party (CCP) are in a race to develop and deploy artificial intelligence. Overwhelming evidence shows that the CCP and its companies are using a suite of escalating actions to steal, subvert, and sabotage American AI. Unless we face this challenge with urgency and resolve, the CCP may succeed in its goal to achieve global AI leadership by 2030 — an unacceptable outcome for American prosperity and power.

Biography

Yusuf Mahmood is Director of AI and Emerging Technology Policy at the America First Policy Institute (AFPI). He holds a J.D. from Harvard Law School and dual bachelor's degrees in economics and philosophy from the University of Maryland, College Park.

[68] National Defense Authorization Act for Fiscal Year 2026, S. 1071, 119th Cong. (2025).