**Testimony before the U.S.-China Economic and Security Review Commission**
**Taking a Bigger Byte: China's Expanding Strategy for Data Dominance**

# Chinese Acquisition of AI Technology

*April 30, 2026*

Dr. Andrew Lohn

I'd like to open by thanking the Commission for this opportunity. It is an honor to be invited to this forum to discuss such an important and rapidly evolving topic.

I lead a small team of researchers at Georgetown University's Center for Security and Emerging Technology. My team's focus is on the intersection of AI and cybersecurity. We also study what is required to create or deploy advanced AI systems in terms of finances, semiconductors, models, and the other underpinnings of the technology. These are all areas of intense competition with China.

This testimony will attempt to place Chinese acquisition of AI technology in the context of their recent history of acquiring technology and sensitive data. It will also outline how those advances in AI might further enhance their cyber capabilities, exacerbating the threat in the future. And I will highlight how the data they stand to steal could become even more impactful as we transition to an AI-enabled society. Finally, I will conclude with some of the actions the U.S. can consider to limit the risks.

## **The Context of Chinese Acquisition**:

More than twenty years ago, the Chinese Titan Rain attacks targeted the U.S. government and defense contractors. Then Operation Aurora pilfered intellectual property from U.S. technology giants in 2009. I remember optimism, such as in the Obama-Xi meeting, that China might agree to not use cyberespionage for commercial gain.[1] The hope was that cyber attacks would be limited to targets that we deem legitimate, excluding corporate intellectual property for the purpose of uplifting Chinese industry. That may or may not have included the subsequent theft of OPM records of government personnel, the Equifax data on about half of all Americans, or the Volt Typhoon incursions into critical infrastructure.[2]

Today I hear less optimism about any ideological restraint and it seems that the CCP views cyber as one tool among many for acquiring U.S. technology. They buy compiled data from data vendors. They use their position in markets such as autonomous vehicles or commercial drones to enhance sensor coverage. They siphon off data in transit around the world as it passes

---

[1] Matt Apuzzo and Michael D. Shear. "U.S. and China Reach Agreement on Economic Espionage." The New York Times, September 25, 2015. https://www.nytimes.com/2015/09/26/world/asia/xi-jinping-white-house.html.

[2] Lily Hay Newman. "All the Ways China's Hackers Stole Data From U.S. Government and Companies." Wired. https://www.wired.com/story/china-equifax-anthem-marriott-opm-hacks-data/.

through their telecommunications infrastructure. And they poach the specialized knowledge of talented individuals such as through the Thousand Talents and related programs.

**Theft of AI**:
Now the technology of the day is AI, and the American companies developing it are often relative upstarts who have little experience as targets of sophisticated espionage campaigns and whose business models often rely on providing access to their technology at extreme scales. These include the familiar AI developers such as OpenAI, Anthropic, and Alphabet, as well as companies that apply AI in innovative ways such as for cybersecurity or military targeting.

Until recently the motivations for China to steal AI have been limited, but that is changing for several reasons. The leading AI models are starting to provide a tactical advantage, especially in cybersecurity. The company XBow topped the bug bounty leaderboard last year, finding more cyber vulnerabilities than any human, and Anthropic's new product Mythos has the cyber world buzzing. As further motivation for theft, Mythos is not available to be tried and tested. If China wants to know what it is capable of, they have to steal it. And beyond the AI systems themselves, these companies have a backlog of thousands of new vulnerabilities, making them even more enticing targets.

Beyond the tactical use for cyber purposes, China may want to steal AI-relevant technology including the model weights, the algorithms or tacit knowledge for training, or the chips they run on. That theft serves a dual purpose of uplifting their domestic industry and subverting U.S. markets and influence. To understand their motivations, we must first understand that China appears to be competing in a different AI race than the United States.

U.S. companies are racing to develop the most capable AI systems regardless of cost. The goal is to develop artificial general intelligence that can replace huge swaths of the workforce and that can improve itself recursively, resulting in runaway advantages that can be turned into profit, security, or wellbeing. The Chinese race on the other hand is to maximize reach. They offer freely-available alternatives that are good enough for many applications despite lower development costs. U.S. companies used to lead in capability among both closed and open models but have lost the lead among open models to Chinese companies. Those low-cost alternatives now undercut the market for U.S. AI services and reduce leverage that the U.S. would otherwise have as the worldwide technology supplier. Even major U.S. companies use Chinese models because the Chinese models can be run on their own servers rather than sending sensitive data to OpenAI, Anthropic, or Alphabet.

It is tempting to mirror our fixation on being at the bleeding edge onto the Chinese but that would be a mistake. U.S. AI technology giants block services from going to China, but China also uses its Great Firewall to block those services from entering China.[3] And despite loosening export controls on U.S. AI chips late last year, Secretary Lutnick informed the Senate last week

[3] "Using ChatGPT in China (Unblocked)." China Survival Kit. https://chinasurvivalkit.com/blog/using-chatgpt-china-unblocked.

that China has not purchased any of those chips yet. Chips are being smuggled illegally into China, but apparently the CCP's policy has been to promote domestic markets even if it means near-term shortages and limitations.

The CCP may also fear reliance on foreign technology that could be cut off in a crisis or conflict. And their authoritarian need to control information makes them acutely aware that AI systems can be imbued with their designer's worldviews or potentially even triggered by its developers. That fear is clear from the requirements that the CCP places on Chinese AI developers.[4] The Chinese are also surely aware of the Crowdstrike report showing that the Chinese DeepSeek model writes more vulnerable code when told that it is working for Tibetans or Uighurs.[5] They probably also read the U.S. studies observing that models can have triggers placed into them with just 250 documents. And while combing over the recent Anthropic code leak this month, they would not have missed the subroutine Anthropic was using to poison models that try to copy their product.[6]

So the Chinese incentives to steal American AI technologies are mixed. Despite banning the models and restricting chip purchases, they are certainly taking active steps to acquire the technology. That includes activities to acquire expertise, hardware, and models. In terms of expertise, the CCP has tempted or coerced AI researchers to China using both carrots such as in the Thousand Talents Program and sticks as in Operation Fox Hunt.[7] Their talent drive also benefited from U.S. policies that repelled foreign talent.[8] In terms of computing hardware, Chinese companies have smuggled billions of dollars of export-controlled AI chips despite restrictions from both the American and Chinese sides.[9] In terms of AI models and systems, all three U.S. AI leaders Alphabet, OpenAI, and Anthropic have accused China of using American outputs to train Chinese copies.[10]

---

[4] International Bar Association. "China Regulation of Artificial Intelligence: Progress and Challenges." IBA Global Insight. https://www.ibanet.org/China-Regulation-of-Artificial-Intelligence-Progress-and-Challenges.

[5] CrowdStrike. "CrowdStrike Researchers Identify Hidden Vulnerabilities in AI-Coded Software." CrowdStrike Blog. https://www.crowdstrike.com/en-us/blog/crowdstrike-researchers-identify-hidden-vulnerabilities-ai-coded-software/.

[6] Anthropic. "Small Samples Can Poison AI Models." Anthropic Research. https://www.anthropic.com/research/small-samples-poison; Alex Kim. "Claude Code Source Leak." March 31, 2026. https://alex000kim.com/posts/2026-03-31-claude-code-source-leak/.

[7] Federal Bureau of Investigation. "FBI Director Christopher Wray's Remarks at Press Conference Regarding China's Operation Fox Hunt." FBI News. https://www.fbi.gov/news/press-releases/fbi-director-christopher-wrays-remarks-at-press-conference-regarding-chinas-operation-fox-hunt.

[8] Remco Zwetsloot et al. "Talent in the Age of AI: How the United States and China Are Competing for Artificial Intelligence Researchers." Proceedings of the National Academy of Sciences 119, no. 49 (2022). https://www.pnas.org/doi/abs/10.1073/pnas.2216248120.

[9] U.S. Department of Justice. "Three Charged with Conspiring to Unlawfully Divert Cutting-Edge U.S. Artificial Intelligence Technology." DOJ Press Release. https://www.justice.gov/opa/pr/three-charged-conspiring-unlawfully-divert-cutting-edge-us-artificial-intelligence.

[10] Bloomberg. "OpenAI Accuses DeepSeek of Using Its Models to Train Competitor." https://assets.bwbx.io/documents/users/iqjWHBFdfxIU/rRmql_jJcxb4/v0/; Google Threat Intelligence.

**AI for Theft**:
Regardless of how much Chinese AI advancement is, or is not, the result of theft, their AI systems are impressively capable, and AI is enhancing cyber attackers. That is happening in two ways that are worth discussing here. One is the ability to make attack tools or exploits. The other is in the process of conducting an attack.

AI is advancing vulnerability discovery as mentioned earlier regarding U.S. companies like XBow and Anthropic's Mythos. Chinese AI systems are also increasingly able to find and exploit vulnerabilities. They are likely a step behind, but it is probably not a big step. Chinese models have tended to be about half a year behind the U.S. leaders. Additionally, if American defenders are not careful, then we will hand adversaries the cyber exploits ourselves. AI is already overwhelming our software maintainers with too many vulnerabilities to address. Some maintainers have closed their bug bounty programs because there is too much to manage. And Anthropic has not released their full vulnerability list because they too are overwhelmed by their own volume.

AI can help write patches for these vulnerabilities, but the hard part is getting those patches into real-world systems.[11] Some of those systems cannot be shut down or restarted easily.[12] And patches occasionally crash the newly updated computers.[13] Those types of concerns can lead to long delays between when a patch is released and when it is taken up across all the systems that need it. In that time, attackers can reverse engineer the patch and attack those who update slowly. That was always a risk, but now AI systems can shorten the attack timeline. An American lead in finding vulnerabilities, and even in developing patches, could inadvertently hand a deluge of intrusion tools to our adversaries.

Even without sophisticated new attack tools, the Chinese are developing AI-enabled cyber operators. We know that because they used Anthropic models to do it. A Chinese attacker strung together about twenty different AI models, each serving different purposes, to conduct a series of attacks against real-world victims.[14] Nothing about the attacks were sophisticated. They used the same boring old tools and techniques, but the AI could try many approaches. Our cyber defenses are built on the principle that no individual defense needs to hold. We layer

---

"Distillation, Experimentation, and Integration: AI Adversarial Use." Google Cloud Blog. https://cloud.google.com/blog/topics/threat-intelligence/distillation-experimentation-integration-ai-adversarial-use; Anthropic. "Detecting and Preventing Distillation Attacks." Anthropic News. https://www.anthropic.com/news/detecting-and-preventing-distillation-attacks.

[11] Andrew Lohn. "Will AI Make Cyber Swords or Shields?" Center for Security and Emerging Technology, Georgetown University. https://cset.georgetown.edu/publication/will-ai-make-cyber-swords-or-shields/.

[12] Tim Starks. "DARPA AI Cyber Challenge Winners Revealed at DEF CON 2025." Cyberscoop. https://cyberscoop.com/darpa-ai-cyber-challenge-winners-def-con-2025/.

[13] Reuters. "Delta Sues CrowdStrike Over Software Update That Prompted Mass Flight Cancellations." October 25, 2024. https://www.reuters.com/legal/delta-sues-crowdstrike-over-software-update-that-prompted-mass-flight-2024-10-25/.

[14] Anthropic. "Disrupting AI Espionage." Anthropic News. https://www.anthropic.com/news/disrupting-AI-espionage.

many defenses and expect that the group will be hard to penetrate. But against an AI adversary that is continually probing and learning where those holes are, that strategy breaks down.[15]

Taken together, there is real concern that AI could substantially boost China's intrusion ability and the number of targets they can prioritize. But so far, that risk is more hypothetical than realized.

On the defensive side, there is hope that AI can help close many of those easy holes. Most of the major intrusions of the past have been the result of poor cyber hygiene. They involve lapses like default passwords, unencrypted datasets, and not implementing multi-factor authentication. Those do not require exceptional creativity to fix. An AI system could scour networks and walk defenders through the steps to harden them. AI may also help accelerate patch uptake, not just patch writing. And AI might help write code that is more secure from the start, reducing the number of vulnerabilities to be discovered.

It is not clear yet who benefits between attackers and defenders. The real-world evidence so far shows offense mostly experimenting, while defenders are starting to be overwhelmed from too much help that could potentially be turned against them.

**The New Targets We Are Creating**:
As AI changes attack and defense, it also changes the targets. Cyber attacks were not a problem before the world became digitized. Now, AI is further digitizing the world both through what it creates and through what it motivates humans to digitize. AI is most useful when it is provided with plenty of information, so people have started building "memory palaces."[16] These are recordings of every interaction, all organized into "rooms" that are connected by digital "halls" and "tunnels." These verbatim conversations cover topics that would have previously only been revealed to a therapist or a spouse, and perhaps not even to them. That data exists in open source memory palaces in formats that are easy to steal and use for malicious purposes, such as coercion.

Companies and governments are restructuring their processes and institutional knowledge to take advantage of AI. China has benefited in the past from stealing blueprints, but it is more helpful to steal the knowledge that turns those blueprints into functioning systems and organizations. Building an extreme ultraviolet lithography machine is a difficult task even if you have stolen designs. Keeping it running when things inevitably go wrong requires decades of experience that have been encoded in the minds of teams of engineers. Agentification across the business world is translating that into digitally-readable formats, increasing the fraction of their organization's secret sauce that can be sucked up by a digital vacuum.

---

[15] Andrew J. Lohn, "Defending Against Intelligent Attackers at Scale," arXiv preprint. "arXiv:2504.18577." https://arxiv.org/abs/2504.18577 (2025).
[16] MemPalace Project. "MemPalace: An Open-Source AI Memory Framework." GitHub. https://github.com/MemPalace/mempalace.

**Policy Interventions**:

Securing our Systems

A natural place to start discussing interventions is at the technical level. We can encourage people and companies everywhere to develop and implement AI for basic cyber hygiene, to use it for patch deployment not just patch writing, and to make software more secure from the start. That could include simple tangible practices such as using AI to write software using less vulnerable languages. It also includes policies and processes such as avoiding the pressure to skimp on the safety or reliability stages that can be a bottleneck to achieving AI's high throughput.

The U.S. should also encourage advanced access to cybersecurity models such as is happening through Anthropic's Project Glasswing.[17] That requires changing some of the norms and principles around vulnerability discovery and disclosure that have been built from decades of hard-won experience. Bug finders typically disclose vulnerabilities to vendors 90 days before disclosing them publicly. Perhaps now, advance notice durations should also account for the number of vulnerabilities disclosed, and more burden may need to fall on the discloser to not only write the patch but to perform more testing or verification to ensure that the patch is reliable enough to be distributed to a wide range of systems.

Mandatory Cyber Standards

From the government agency side, a lot of the burden falls on the Critical Infrastructure and Security Agency (CISA) which is understaffed from layoffs, has had leadership uncertainties, and has been plagued by intermittent shutdowns as part of the Department of Homeland Security.

Much of the burden also needs to fall to the technology developers themselves. As companies develop increasingly critical technology, it is reasonable to expect them to implement correspondingly rigorous defenses. There have been few bills and fewer laws driving mandatory cybersecurity measures among these companies. SB-53 in California and the RAISE Act in New York make generic demands for incident reporting and abstract cyber defenses.

Given the growing risks, it is worth establishing a federal stance with more specificity, and it can be done without imposing excessive regulatory burden on companies that cannot afford it. SB-53 defines covered organizations in two ways. One refers to the computing budget used on the expense side, and the other refers to at least $500 million on the revenue side. Both assure that covered entities have enough budget that some spending on cybersecurity and compliance would be warranted and not overly burdensome.

If federal law is infeasible, another approach is to follow the example of pipelines and rail. The Transportation Security Agency, which is the Sector Risk Management Agency for those components of critical infrastructure, enacted and upheld cybersecurity directives in both the Biden and Trump administrations. As the Sector Risk Management Agency for Information

[17] Anthropic. "Project Glasswing: Advanced Access to Cybersecurity Models." https://www.anthropic.com/glasswing.

Technology, it may be possible for CISA to direct AI model providers to enact mandatory cybersecurity measures or at least to report breaches.

Distillation

The policies and techniques described so far relate to traditional cybersecurity operations but are less focused on adversarial distillation where the normal operation of the models is used to extract their knowledge and transfer it to other developers. It is not clear how to prohibit adversarial distillation. It is not certain whether it is prohibited by the Computer Fraud and Abuse Act related to cyber attacks. It is not certain whether it is prohibited by the Digital Millennium Copyright Act. And it is not certain whether it amounts to Economic Espionage according to Title 18. Two weeks ago, the Senate introduced the Deterring American AI Model Theft Act which might not have the teeth to fully live up to its title, but might help move the needle on making these techniques explicitly illicit and providing some measure of response.

Both distillation and cyber threats are growing against all of the major AI developers, and they could benefit from government coordination for information sharing, both among each other and with the intelligence community. Some of that coordination can be accomplished through government agencies such as CISA or the Center for Standards and Innovation. The NSA's AI Security Center and Cybersecurity Collaboration Center can also help identify threats to inform companies and receive threat intelligence from companies to inform government operations and to inform other likely targets across industry. Stabilizing and staffing CISA and the NSA information sharing organizations should be a priority if the goal is to prevent China's cyber teams from exfiltrating American technology and personal data.

Hardware

As for the hardware that is being funneled to China, the Bureau of Industry and Security, the organization responsible for tracking and recovering these infractions, has been underfunded. Preventing even a small fraction of the billions of dollars of illicit shipments would justify a boost in their funding. There will also be future bills similar to the one put forward last year by Representatives Foster and Huizenga to collect location information on these chips. Location verification has downsides, and a forthcoming report from our team at CSET will help navigate benefits and shortcomings, but the monetary and practical value of these chips is very high. Some sacrifices are likely to be worthwhile, even if the regulation places some burden on American chip designers and even if there are ways to circumvent the controls.

Talent

As for talent, the United States should enact policies and laws that help attract and retain foreign talent such as through various visa programs and by admitting foreign students. The United States has a good record of retaining those talented individuals while creating a brain drain abroad. As part of that effort, the United States should work to create a climate that is welcoming to foreigners and undercut China's efforts to pull those individuals back to China.

Security Assurances

Finally and most broadly, the United States, and Congress in particular, should be cautious with the sacrifices it makes in pursuit of AI, given the alerts raised by the American AI labs that China is freeloading off American investment. If America is to reallocate energy and water infrastructure, to finance or backstop corporations and datacenter buildouts, and to provide regulatory relief for construction, privacy concerns, copyright infringement, and psychosocial harms, the benefits should not accrue to China. Congress, and the American taxpayer, should demand assurances that AI developers can protect the technology as a precondition to receiving our support.