


JUNE 2026

Red Lines

Understanding the National Security Risks of China's Advanced AI

Daniel Remler



About the Author



Daniel Remler is a senior fellow with the Technology and National Security Program at the Center for a New American Security (CNAS). His research focuses on the implications of artificial intelligence (AI) and emerging technologies for U.S. national security, foreign policy, and strategic competition, including risks from advanced

AI systems. Prior to joining CNAS, Remler served as a policy advisor in the Office of the Special Envoy for Critical and Emerging Technology at the U.S. Department of State, where he led AI policy and cowrote the department's first technology diplomacy strategy. He previously worked as a journalist with *The Economist*, where he covered American politics and public policy. Remler holds an MPA from the Harvard Kennedy School and a BA in economics and history from the University of California, Berkeley.

About the Technology and National Security Program

The CNAS Technology and National Security Program produces cutting-edge policy research to secure America's edge in emerging technologies, while managing potential risks to security and democratic values. The program produces bold, actionable recommendations to drive U.S. and allied leadership in responsible technology innovation, adoption, and governance. The Technology and National Security Program focuses on three high-impact technology areas: AI, biotechnology, and quantum information sciences. It also conducts cross-cutting research to strengthen U.S. technology statecraft to promote secure, resilient, and rights-respecting digital infrastructure and ecosystems abroad. A focus of the program is convening the technology and policy communities to bridge gaps and develop solutions.

Acknowledgments

The author thanks Seth Center, Vivek Chilukuri, Janet Egan, David Lin, Simon Nash, and Paul Scharre for their valuable feedback and suggestions on earlier drafts of this report. He is also grateful to the dozens of experts in government, industry, and civil society who participated in a roundtable at CNAS or who agreed to be interviewed as part of this research project. The report would not have been possible without the research support of Benjamin Hayum, as well as the editorial

and design contributions of CNAS colleagues Maura McCarthy, Melody Cook, Caroline Steel, and Emma Swislow. This report was made possible with the generous support of Coefficient Giving.

As a research and policy institution committed to the highest standards of organizational, intellectual, and personal integrity, CNAS maintains strict intellectual independence and sole editorial direction and control over its ideas, projects, publications, events, and other research activities. CNAS does not take institutional positions on policy issues, and the content of CNAS publications reflects the views of their authors alone. In keeping with its mission and values, CNAS does not engage in lobbying activity and complies fully with all applicable federal, state, and local laws. CNAS will not engage in any representational activities or advocacy on behalf of any entities or interests and, to the extent that the Center accepts funding from non-U.S. sources, its activities will be limited to bona fide scholastic, academic, and research-related activities, consistent with applicable federal law. The Center publicly acknowledges on its [website](#) annually all donors who contribute.

Table of Contents

| | |
|----|---|
| 01 | Executive Summary |
| 02 | Introduction |
| 03 | I. How China's Advanced AI Threatens U.S. National Security |
| 06 | II. Risk Assessment: The Current State of Chinese Advanced AI Systems |
| 16 | III. Conclusion |



Executive Summary

CHINESE ADVANCED ARTIFICIAL INTELLIGENCE (AI) systems pose a serious and growing threat to U.S. national security. At least seven Chinese developers now produce systems with formidable capabilities across coding, reasoning, multimodal recognition, and agentic tasks—systems that are released with open weights, offered via application programming interface (API) at prices designed to undercut American competitors, and available for download by anyone in the world. The Chinese Communist Party (CCP), which collapses the boundaries between state, military, and private sector, treats these systems as instruments of political control, economic dominance, and great-power competition.¹ Some of these risks are inherent to any sufficiently capable AI system that lacks adequate safeguards. But the most significant risks are products of the political system that builds, shapes, and deploys these systems. This party-state does not tolerate independent power centers and treats AI as a tool of statecraft across every dimension of strategic competition.²

This report proposes a framework for understanding these risks across three domains and two vectors. In the **kinetic domain**, Chinese AI systems enhance military capabilities and offensive cyber operations and raise concerns about biological weapons development. In the **cognitive domain**, they enable more effective censorship, surveillance, influence campaigns, and espionage. In the **economic-technological domain**, they drive industrial dominance and create dependencies that extend China's reach in emerging and advanced economies alike. These risks affect the United States through two vectors: **state instrumentalization**, in which the CCP directly wields AI systems, and **proliferation and dependency**, in which Chinese systems spread globally through open-weight release, model compression, and aggressive pricing. Widespread adoption of Chinese AI systems creates structural dependencies that give the CCP leverage, embed CCP ideology and security vulnerabilities into foreign systems, and expand the attack surface available to Chinese intelligence—even when no one in Beijing lifts a finger.³

A detailed assessment of the current capabilities, design choices, and security vulnerabilities of systems from China's seven leading AI developers (Alibaba, Baidu, DeepSeek, MiniMax, Moonshot, Tencent, and Zhipu) reveals threats that are concrete and, in several cases, immediate. Chinese systems can already contribute meaningfully to offensive cyber operations. Ideological alignment with the CCP is deepening with every new model. DeepSeek-based agents are 12 times more likely to follow malicious instructions than their U.S. counterparts.⁴ These findings underscore why this report argues for understanding risk in absolute terms, focusing on what Chinese systems can do, not just how far behind they are.

The report offers six policy recommendations:

1

The Department of Commerce should publish national security risk assessments of Chinese advanced AI systems no more than 72 hours after their release.

2

The Cybersecurity and Infrastructure Security Agency should issue cybersecurity alerts and advisories on Chinese advanced AI systems and establish the AI Information Sharing and Analysis Center called for in the AI Action Plan.

3

The Department of Commerce should establish security testing best practices for cloud service providers that deploy and monitor AI models developed by untrusted third parties.

4

The Department of Energy should use the Genesis Mission to establish a classified adversarial testing program for Chinese advanced AI systems.

5

The Department of State should convene monthly meetings with a core group of U.S. allies to share information on China's AI ecosystem and coordinate collective action.

6

The Department of Commerce should publish semiannual reports on the state of China's AI ecosystem to inform policy action and congressional oversight.

Introduction

CHINA'S SPRING FESTIVAL traditionally marks a time for renewal—to reconnect with family, present one's best self to the world, and look forward to the year. The holiday has accordingly become a showcase of China's artificial intelligence (AI) power. In the weeks surrounding Lunar New Year 2026, Chinese AI developers released a volley of new systems that demonstrated the breadth and strategic ambition of the country's AI ecosystem. The February 11 release of Zhipu's GLM-5 kicked off the holiday season with a new 744-billion-parameter system compatible with Huawei chips that scored just behind GPT-5.2 and Claude Opus 4.5 on popular coding benchmarks.⁵ Moonshot, ByteDance, and Alibaba all released upgraded systems of their own over the following weeks.⁶

Model releases were only part of the story. Each system was released with open weights under permissive licenses. Each was offered via application programming interface (API) at a fraction of the cost of its American equivalent. And each was available for download by anyone in the world. Tencent, ByteDance, Alibaba, and Baidu used the same holiday to run an AI adoption war, embedding their chatbots into red envelope giveaways, the China Central Television Spring Festival Gala, and e-commerce workflows. Tencent's Yuanbao chatbot passed 50 million daily active users during the holiday alone.⁷ The frenzy intensified in March, when OpenClaw, an open-source agent framework, went viral across China as millions of users connected it to Chinese models such as Qwen to automate everyday tasks—prompting every major Chinese technology company to launch its own version within days.⁸

Behind the symbolism of releasing new models around the Spring Festival lies a dynamic ecosystem that now fields at least seven Chinese developers producing systems that pose serious national security challenges for the United States. The Chinese Communist Party (CCP) treats the advanced AI emerging from this ecosystem as essential to its military modernization, political control, economic ambitions, and global influence. The recently released *15th Five-Year Plan* reinforces this priority, listing AI as the top “frontier technology” ahead of quantum computing, fusion energy, and other fields.⁹ Beyond indigenous innovation, the CCP seeks advantage across every segment of the AI supply chain—in software through adversarial distillation (training on the outputs of superior U.S. models to replicate their capabilities at a fraction of the cost), in hardware through access to banned chips, and in integration by privileging partnerships between Chinese firms.

The Trump administration's July 2025 AI Action Plan included some initial direction to grip the challenge of Chinese AI.¹⁰ It called on the U.S. government to evaluate national security risks in frontier Chinese models, strengthen the enforcement of AI compute export controls, align protection measures with allies,

and promote mature federal capacity for AI incident response. But execution has not kept pace with ambition. As of publication, the Department of Commerce has published only three assessments of Chinese AI models, both delayed weeks or months after the models were released and the technical analyses had been completed. No Chinese AI developer has been subject to U.S. government sanctions of any kind since 2024. Allied coordination remains limited and ad hoc. Meanwhile, the direct threat to Americans has escalated. Chinese AI systems are powering cyber campaigns against U.S. infrastructure, embedding CCP ideology into software adopted by millions, and creating espionage vulnerabilities that Chinese security services are uniquely positioned to exploit.¹¹

In September 2025, Director of the Office of Science and Technology Policy Michael Kratsios warned the UN Security Council that “the improper use of AI systems can erode deterrence, create destabilizing effects, and reinforce systems of political control and social engineering.”¹² China's advanced AI systems do all three, posing risks to national security across three domains: kinetic, where they enhance military and cyber capabilities; cognitive, where they enable censorship, surveillance, and espionage; and economic-technological, where they drive industrial dominance, threaten to dominate commercial platforms, and create dependencies.

The result is an analytical vacuum at the center of the AI policies of the United States and its allies.

These risks are not incidental features of model design. They flow from the nature of the CCP, a party that collapses the boundaries between the party and society and that views AI as an instrument of political control, economic dominance, and great-power competition. These risks reach the United States and its allies through two vectors: direct state instrumentalization by the CCP, and proliferation and dependency as Chinese systems spread globally through open-weight release, aggressive pricing, and integration into critical infrastructure.

The lack of clarity over these risks has been fed in part by the Western AI community's fixation on U.S. frontier systems. The most rigorous research and most consequential regulatory interventions have focused almost exclusively on American developers, inflating the risks of technology built under the rule of law while giving China's developers a free pass. The result is an analytical vacuum at the center of the AI policies of the United States and its allies.

The report makes the case for measuring these risks in absolute terms based on what Chinese systems can already do, not just how far behind they trail their U.S. peers. The prevailing “AI race” framing obscures whether Chinese systems have crossed specific capability thresholds that endanger American

security. This report assesses the capabilities, vulnerabilities, and design of AI systems from China's seven leading developers, pegging them to actual use cases where feasible. A Chinese AI system that can conduct an autonomous cyber campaign, suppress dissent at scale, or lock developing countries into Chinese infrastructure threatens U.S. national security regardless of whether an American system is better. Policy built on relative rankings breeds complacency when the U.S. lead grows and defeatism when the gap narrows. Policy built on understanding absolute capabilities produces specific objectives and sharper tools.

Section I of this report situates AI risks to U.S. national security within the CCP's broader strategic logic, providing a framework for understanding these risks. Section II assesses the current capabilities, design choices, and security vulnerabilities of advanced AI systems from China's seven leading AI developers across this framework. The recommendations subsection of Section III offers six policy prescriptions, most of which are designed to close analytical gaps that unclassified and classified testing, allied coordination, and public reporting could address, but that the U.S. government has not yet prioritized.

How China's Advanced AI Threatens U.S. National Security

This section proceeds in three parts. It first examines how the CCP's core interests and its legal, regulatory, and institutional apparatus make advanced AI central to party strategy. It then introduces a framework for categorizing the resulting risks to national security across three domains—kinetic, cognitive, and economic-technological—and across two vectors of state instrumentalization and of proliferation and dependency.

AI and the Party's Core Interests

AI serves the CCP as both an object of its core interests and as a tool to employ in the pursuit of these interests. The CCP has built a comprehensive apparatus—legal compulsion, strategic planning, regulatory control, military-civil integration, and international standard-setting—to ensure technology development serves the party at home and extends its influence abroad.¹³ The party defines its core interests—sovereignty, security, and development—as nonnegotiable and unlimited in scope.¹⁴ Coupled with a Leninist system that tolerates no autonomous power centers, these interests collapse boundaries between domestic and foreign policy, and between the party, the state, and the private sector.¹⁵ The sections that follow explain why this dynamic makes AI central to CCP strategy, how Beijing has organized to exploit AI technology, and what risks it creates for U.S. national security.

The CCP's core interests are inescapably global. Asserting sovereignty over Taiwan requires the capability to deter or defeat American intervention, while securing state control over the Chinese internet requires shaping global internet governance.¹⁶ Protecting the party's security means preventing external actors from undermining its legitimacy through political, economic, or military means, a nonnegotiable imperative that cannot coexist with a U.S.-led international order. Sustaining development means achieving the technological self-reliance that General Secretary Xi Jinping has elevated to a strategic priority, insulating China from American pressure.

Beijing has codified these obligations to support state intelligence and security operations into law.¹⁷ The *2017 National Intelligence Law* compels all organizations and citizens to “support, assist, and cooperate with national intelligence efforts.”¹⁸ The *2017 National Cybersecurity Law* mandates cooperation with the security services.¹⁹ In China's political-economic system, any company can function as an instrument of state power. These authorities provide the legal backbone for the CCP's strategy to develop and use AI in all its dimensions.

The party's conviction that technology is the fulcrum of great-power competition dates to at least 1978 and Deng Xiaoping's

“Four Modernizations.”²⁰ It has only intensified since. Xi frames technology-driven “self-reliance” and “self-strengthening” as essential to China’s security, describing technology as “the main arena of international competition.”²¹

In China’s political-economic system, any company can function as an instrument of state power.

As a general-purpose technology, AI can be applied to every dimension of the party’s core interests.²² AI-enabled surveillance makes authoritarian control of 1.4 billion people possible at a scale no human bureaucracy could match.²³ Capabilities that power predictive policing and content moderation at scale also drive the People’s Liberation Army’s (PLA’s) vision of “intelligitized warfare,” compressing military decision-making that could shift the balance in a Taiwan contingency.²⁴ Economic growth, military modernization, and surveillance infrastructure all draw on the same well of AI research, talent, and compute—and a breakthrough in any domain strengthens the others.

AI is also distinct as a cultural technology. Language models encode values, ideological assumptions, and cultural knowledge. Systems trained on Western data reflect liberal assumptions about free expression and individual autonomy.²⁵ In 2023, Chinese military commentators argued that Western AI systems are trained on data reflecting Western political values and could subtly influence users’ behavior—a concern that extends the party’s long-standing anxiety about ideological infiltration codified in a 2013 directive to guard against “universal values” and “Western constitutional democracy.”²⁶ For a Leninist organization premised on ideological control, ceding AI development to foreign or domestic actors outside the party risks embedding dissent into China’s technological infrastructure.²⁷

The CCP’s response is to build its own AI capabilities while shaping the technology’s development and global diffusion. The military-civil fusion strategy, elevated to a national priority in 2017, aims to eliminate barriers between civilian innovation and military applications, ensuring that commercial AI breakthroughs are available to the PLA.²⁸ Practical friction slows adoption, but no Chinese company has the legal basis or political space to refuse when the party asks. The 2017 *New Generation AI Development Plan* complemented this strategy by directing state resources toward specific AI capabilities, designating national champions and setting ambitious targets while explicitly mandating two-way transfer between military and civilian AI research.²⁹ More recently, the People’s Republic of China’s (PRC’s) State Council’s 2025 AI Plus initiative, now enshrined as a priority in the 15th *Five-Year Plan*, aims to drive adoption across the government and the economy.³⁰ The Digital

Silk Road exports Chinese AI and surveillance systems to overseas markets.³¹

Beijing’s AI governance framework reinforces the party’s grip over the design and use of Chinese AI systems. The 2022 *Algorithmic Recommendation Provisions* and 2023 *Generative AI Measures* mandate that AI systems “reflect socialist core values,” refuse to generate content that “subverts state power,” and give authorities access to training data and model parameters.³² The 2025 *AI Safety Governance Framework 2.0* included “challenges to existing social order” and “supply chain safety” in its list of “safety risks.”³³ The Cyberspace Administration of China maintains a central registry of all public-facing AI products, requiring developers to provide documentation and pass a mandatory, broadly defined security assessment before deployment.³⁴ Technical standards translate these directives into mandatory tests.³⁵ China’s AI governance framework is not about safety in the Western sense—it is an instrument of political control, designed to ensure AI systems reinforce party authority and remain visible, interpretable, and subject to oversight by the state.

That a party-state so intolerant of independent power centers is nonetheless eager to diffuse advanced AI at home and abroad is testimony to the robustness of this apparatus of control. Diffusion extends the party’s reach rather than diminishes it—and the dangers travel with the technology, embedding party ideology, introducing exploitable vulnerabilities, and creating architectural dependencies that follow Chinese systems into foreign infrastructure.

Indeed, the party’s own direct efforts to shape AI development do not stop at China’s borders. Beijing has pursued a long-standing campaign to elevate the United Nations as the central venue for AI governance, where it can build coalitions among developing countries and exclude nongovernmental organizations.³⁶ The 2023 *Global AI Governance Initiative* and 2025 *Action Plan* position state sovereignty as the organizing principle of global AI governance, the international analog of the party’s domestic control.³⁷ Terms such as “AI safety” left undefined for international consumption mask the CCP’s authoritarian conception at home.³⁸ At the same time, national guidelines emphasize shaping international technical standards, setting a target of “at least 20” standards.³⁹ This strategy mirrors the domestic playbook, promoting high-level norms that enshrine state control and technical standards that entrench Chinese technology.

These dynamics produce specific, measurable risks to national security. The framework in the next part of this section is intended to categorize these risks, and the assessment in Section II documents them.

Framework for Categorizing AI Risks to U.S. National Security

U.S. policymakers have typically developed policy to address Chinese AI risks through disconnected silos. Export controls have been justified on military and intelligence grounds.⁴⁰ Programs to promote U.S. AI technology abroad focus on commercial competition with Chinese firms in third countries to address economic dependencies, but remain disconnected from other dimensions of the challenge from Chinese AI.⁴¹ The AI safety community has focused primarily on misuse and loss of control as technical problems, an approach that has produced valuable research but that has ignored geopolitical imperatives. Some in the community have advocated cooperation with Beijing on shared risks—a position that, whatever its merits on narrow technical questions, is willfully ignorant of the CCP's treatment of AI safety as an instrument of state power. These communities rarely engage one another. The result is policy incoherence.

An integrated framework can connect these strands. Any categorization simplifies, but policymaking requires heuristics, including ones that map onto how governments organize. This report argues for understanding the risks from Chinese AI through two vectors and three domains. Vectors capture how the party pursues its interests, either through direct instrumentalization of AI by state actors or through global proliferation that extends influence and creates dependencies. Domains describe the types of strategic capability—kinetic, cognitive, and economic-technological—that serve CCP interests.

Some capabilities cut across domains but can be disaggregated by use case. Cyber operations are a clear example. This report addresses cyber offense under the kinetic domain and cyber-enabled espionage and intelligence collection under the cognitive domain, reflecting the primary purpose each serves, while recognizing that a single operation can advance CCP interests across all three.

This framework focuses on risks that flow from CCP rule. Some risks from advanced AI, such as uncontrolled autonomous behavior, may not be unique to Chinese systems and are addressed by a separate body of research. This report engages them where they intersect with CCP-specific dynamics documented here, particularly where minimal safeguards amplify risks.

Absolute improvements in China's AI capabilities compound risk across all six categories. A more capable AI model can enhance military targeting, scale influence operations, and accelerate industrial catch-up while creating more attractive technology for global adoption. No individual risk can be assessed in isolation from the trajectory of China's AI ecosystem. What matters is what Chinese systems can do to threaten U.S. national security, not how far behind they are.

The framework is also useful for evaluating policy. Technical solutions to detect synthetic content address one aspect of “information control.” The U.S. AI Exports Program targets “systemic presence” and “economic leverage” by promoting U.S. technology in third countries. Not every response must operate at the model level, and effective measures may be better targeted at downstream use cases or upstream inputs. But the most effective policies degrade China's overall capacity to develop and use AI, cutting across every risk category at once.

The scale and interconnection of these risks demand that policymakers have far better understanding of the Chinese AI ecosystem—its development, capabilities, and trajectory—than they do at present. The next section provides this assessment.

Table 1 | Categories of National Security Risks from Chinese Advanced AI Systems

Vectors define the mechanisms of risk transmission: *state instrumentalization*, in which the Chinese Communist Party (CCP) directly leverages AI systems to advance its strategic objectives, and *proliferation and dependency*, in which the global diffusion of Chinese AI creates structural leverage.

Domains classify the strategic functions of the capability: *kinetic* (physical force projection), *cognitive* (information manipulation and influence), and *economic-technological* (industrial competitiveness and supply chain control).

| Domains | Vectors | |
|---|--|--|
| | State Instrumentalization <i>How the CCP directly employs AI capabilities to advance its core interests</i> | Proliferation & Dependency <i>How the global spread of Chinese AI technology structurally advances the CCP's core interests</i> |
| Kinetic <i>Physical force or the threat thereof</i> | Military modernization Integration of AI into weapon systems, command networks, and intelligence analysis to enhance combat capabilities, operational planning, and rapid decision-making, including in cyber operations | Proxy empowerment Open-weight release and weak guardrails structurally enable access to Chinese AI for use by adversarial states and nonstate actors for military, cyber, biological, and other dangerous applications |
| Cognitive <i>Belief and perception formation</i> | Information control AI enables automated mass surveillance, censorship, and influence operations at scale to maintain domestic control and shape foreign perceptions | Systemic presence Deployed systems create vectors for intelligence collection, data access, and propagation of CCP-aligned concepts |
| Economic-Technological <i>Control over resources, technologies, networks, and markets</i> | Industrial dominance AI advances industrial policy goals, enabling technological self-sufficiency and competitive advantages in strategic sectors | Economic leverage Integration in critical infrastructure, government functions, and business operations creates dependencies and opportunities for coercion and market dominance |

Risk Assessment: Current State of Chinese Advanced AI Systems

ADVANCED AI SYSTEMS OFFER the clearest window into China’s AI ecosystem and the risks it poses to U.S. national security. These systems serve as tools for the CCP and barometers for the development of China’s wider AI ecosystem. This section first identifies the seven Chinese developers whose systems warrant the closest scrutiny, then assesses those systems across the three domains identified in Section I—kinetic, cognitive, and economic-technological.

What makes advanced AI systems useful as tools and barometers is that they are the most general-purpose form of AI. DeepSeek-R1, for example, was quickly integrated by the PLA, industrial firms, and surveillance operators alike.⁴⁴ Open-weight release accelerates diffusion, allowing any actor to download and modify these models to suit its own needs.⁴⁵ Improvements in these systems cascade into narrower applications built on top of them, meaning a better Chinese model can be deployed across all six risk categories identified in Section I. A model-level focus reframes the threat from Chinese digital technology—where Chinese systems embedded in foreign code, workflows, and decisions now matter more than where Chinese hardware and networks are built. Local deployment does not contain these risks.

One counterargument is that it can—locally hosted open-weight models eliminate acute risks. Data does not flow to China and API-level manipulation is impossible. Indeed, not all uses of Chinese AI are equally dangerous. But widespread adoption creates structural dependencies, normalizes Chinese AI in enterprises, and embeds persistent characteristics. A local DeepSeek instance still carries CCP-aligned ideology and elevated prompt injection vulnerability. The line between critical and noncritical use is also unstable, as code written for a start-up today may end up with a defense contractor tomorrow, while widespread adoption can create leverage beyond software. As models such as Qwen become the default for on-device AI, their developers gain influence over chip design by sharing or withholding optimization information with hardware manufacturers.⁴⁶

Tracking China’s advanced AI systems can also help calibrate U.S. policy. For example, export controls on advanced semiconductors and manufacturing equipment have, as intended, limited the abilities of Chinese firms to train models and run inference across the economy.⁴⁷ Developers such as Alibaba and DeepSeek have acknowledged the damage of chip shortages, with Zhipu Cofounder Tang Jie going so far as to admit in early 2026 that “the truth may be that the gap is actually widening,” while his own company says it is “pushing every chip to its limit just to serve inference.”⁴⁸ Model capabilities and adoption are therefore

Definition of Advanced AI Systems

This report uses “advanced AI systems” as a catch-all term to describe AI models that are often called generative AI, large language models (LLMs), or foundation models. Advanced AI systems also encompass other commonly used tools, such as chat interfaces and agent harnesses (software that enables AI models to take actions, such as web browsing, executing code, etc.).⁴² This term is used widely, if inconsistently, in international forums, and avoids confusion with emerging regulatory definitions.⁴³

a proxy for the combined effect of U.S. export controls and China’s countervailing policies, above all adversarial distillation, which allows Chinese developers to close capability gaps without the compute needed to build comparable models independently. Analysis by the Institute for Progress estimates that without advanced chip exports, U.S. compute capacity exceeds China’s by more than tenfold—a gap that shapes the capability assessed in this report.⁴⁹

Understanding advanced AI systems requires understanding their developers. The caricature that Chinese developers focus on applications and eschew the American pursuit of ever-larger systems overstates reality.⁵⁰ Chinese developers produce large flagship models as well as systems of different sizes and specializations. Alibaba, like Google, continues to scale its Qwen models while integrating them across its platform.⁵¹ Zhipu, like Anthropic, has built capabilities directly into its AI system that other developers offer as stand-alone tools.⁵² In a fiercely competitive market, China’s leading developers are building genuinely sophisticated systems.⁵³ They are the vanguard of the ecosystem and the focus of this assessment.

Defining the Chinese AI Vanguard

A combination of measures can identify key players in the Chinese ecosystem. Three categories, each with its own metrics, converge on the most advanced AI systems in China. These include intelligence, real-world value, and technical design. Taken together, they distinguish theoretical capabilities from realized value, and model performance from the technical resources and design behind it (see Table 2). A developer need not lead in every category to be strategically significant. What distinguishes the developers identified here is competitive performance across enough categories to sustain a dynamic AI ecosystem. These categories synthesize metrics used across the most widely cited AI evaluation platforms.

These measures identify seven Chinese developers as ecosystem leaders: Alibaba, ByteDance, DeepSeek, MiniMax, Moonshot, Tencent, and Zhipu.⁵⁴ Their flagship systems rank ahead of Chinese peers but lag behind U.S. models from Anthropic, Google, and OpenAI.⁵⁵

Table 2 | Metrics of Model Capabilities

| Category | Subcategory |
|------------------------------|---|
| Knowledge/ Intelligence | Amortized knowledge: How well a model can solve a problem using internal knowledge that does not require additional computational resources |
| | Reasoning/agent intelligence: How well a model can solve a problem by employing computational resources to “reason” to a solution |
| | Time horizon: How quickly a model can solve a problem quantified in terms of human-equivalent worker-hours |
| | Multimodality: How well a model can solve problems across different input modalities (e.g., text, image, video, audio) |
| Real-World Value | Economic value: How well a model performs at tasks particular to a specific profession or industry |
| | Research value: How many tasks within various research domains an AI model can perform |
| | Physical use value: How well a model can perform on tasks with real-world effects, including military operations, cyberattacks, robotics, etc. |
| Technical Design and Compute | Compute quantity: How much compute is used in training the model, measured in floating point operations per second (FLOPs), gigawatts, etc. |
| | Scaling efficiency: How efficient the model is at learning per unit of data or FLOPs spent in training |
| | Inference cost: How much it costs to deploy the model at scale (e.g., cost-per-token offered via API) |
| | Model tooling/orchestration: How effective the tool scaffolding is and integrations around the model (e.g., an agent harness) |

These developers capture the technical sophistication and diversity of the ecosystem, underlining their importance as barometers of China’s AI prowess. Alibaba, ByteDance, and Tencent leverage AI across sprawling platforms—Alibaba in e-commerce and cloud computing, ByteDance in social media, and Tencent in messaging and gaming.⁵⁶ DeepSeek operates as a research project for its parent hedge fund, with chief executive officer Liang Wenfeng positioning it as a pure research lab dedicated to achieving artificial general intelligence.⁵⁷ Moonshot and Zhipu, both Tsinghua University spinouts, have a similar vision to DeepSeek.⁵⁸ Zhipu derives much of its revenue from state-owned enterprises and has close ties to the government, enough to land it on the U.S. Department of Commerce Entity List in January 2025.⁵⁹ Along with MiniMax, it recently debuted on the Hong Kong stock exchange in pursuit of needed capital and customers, while Moonshot remains private.⁶⁰

The rest of this section assesses how these developers’ systems serve CCP interests across the three domains identified in Section I. This assessment draws on public benchmarks, model cards, threat reports, and reported uses. Chinese AI developers are opaque, Western scrutiny has been limited, and some information is classified by necessity, meaning analytical gaps persist.

Where available, this assessment uses U.S. systems as reference points, not to frame the analysis in relative terms, but because independent evaluations of Chinese systems are scarce and U.S. models provide a familiar baseline for what capability levels mean in practice.⁶¹ Closing this evaluation gap is itself a focus of Section III.

The Kinetic Domain

Advanced AI systems provide the PLA with one pathway to “intelligentized warfare,” the operational concept at the heart of its modernization strategy.⁶² Off-the-shelf systems are not purpose-built for direct military use, but Chinese systems possess demonstrable capabilities that, whether used by the CCP or other malign actors, put Americans at risk.

This section examines three areas in which Chinese AI capabilities could pose kinetic risks: conventional military applications, including the integration of AI into PLA planning, logistics, and command systems; offensive cyber operations; and biological weapons, in which capable models and weak safeguards may create proliferation risks.

Conventional Military-Relevant Capabilities

Advanced AI systems can enable militaries to act at greater scale, speed, and coordination.⁶³ While the readiness of Chinese AI systems for military applications cannot be fully assessed through unclassified sources alone, public benchmarks and reporting from both the United States and China are suggestive of how these tools are now and may soon be used. Agentic and software engineering capabilities, coupled with multimodal recognition and multistep reasoning, could prove highly effective in command and control, logistics, decision support, and intelligence analysis.⁶⁴

The most immediate applications are in the higher-latency, analysis-heavy tasks preceding and following kinetic engagements, such as synthesizing intelligence streams, generating and prioritizing target packages, tracking logistics, and assessing battle damage.⁶⁵ The U.S. military’s early experience in Iran in 2026, where Anthropic’s Claude models reportedly assisted in targeting that enabled strikes on more than 1,000 targets in the first 24 hours, illustrates how quickly these capabilities translate into operational advantage.⁶⁶

Evidence suggests the PLA is experimenting with these uses.⁶⁷ PLA procurement data analyzed by the Center for Security and Emerging Technology (CSET) reveals increasing adoption of Chinese advanced AI systems, particularly from Alibaba and DeepSeek, with at least some military elements beginning to apply them to intelligence tasks.⁶⁸ This suggests China’s AI developers are increasingly implicated in military-civil fusion, as nontraditional vendors without self-reported state ownership represented 70 percent of all entities awarded AI-related contracts from 2023 to 2024 in the CSET dataset.⁶⁹ Beihang

University, a key PLA research partner, has filed patents using DeepSeek to improve drone swarm decision-making for “low, slow, small” aerial threats.⁷⁰ PLA strategists continue to explore the uses and limits of AI in the military, echoing U.S. debates.⁷¹

Two leading benchmarks of software engineering proficiency suggest Chinese systems are already useful for the PLA. On SWE-Bench Pro, which tests models’ abilities to resolve complicated software engineering issues, leading Chinese systems achieve about 50 percent accuracy, compared to more than 55 percent for GPT-5.2 and Claude Opus 4.6.⁷² Terminal-Bench 2.0, which tests AI agents on how well they can autonomously resolve software engineering issues through the command line, scores Chinese models Kimi K2.5 and GLM-5 at 50.8 percent and 56.2 percent, respectively, compared to 54 percent and 65 percent, respectively, for GPT-5.2 and Claude Opus 4.6.⁷³

These scores are good enough for selective adoption in the intelligence analysis and planning functions in which AI has proven most valuable, but they remain immature for wider use.⁷⁴ The energy needed for constrained onboard hardware is prohibitive, and the velocity of weapon systems is fundamentally incompatible with the latency of networked model inference.⁷⁵ The battlefield environment, with sensor degradation, jamming, and adversarial conditions, is an environment in which current AI systems do not perform well.⁷⁶ Even at the operational level, integration with legacy software and data systems is a significant constraint.⁷⁷ The organizational and doctrinal challenges of embedding AI into existing command structures—establishing trust, managing failure modes, and maintaining accountability—will slow adoption regardless of raw capability scores.⁷⁸ Chinese AI may not yet be enabling battlefield weapons, but these constraints are engineering and bureaucratic problems, not fundamental limits.

As capabilities improve, model compression helps overcome technical limitations. Quantization shrinks model weights by reducing memory footprint and power consumption with minimal accuracy loss.⁷⁹ Distillation—here, the legitimate practice of training a smaller model on a developer’s own larger system—produces efficient variants that preserve most of a flagship model’s capability, particularly for narrow task categories, at a fraction of the computational cost.⁸⁰ Modern weapon platforms lack the hardware to run large models, but distilled variants from Alibaba or DeepSeek could plausibly run on operational-level hardware comparable to that of the U.S. military.⁸¹ Distillation and onboard compute may bring frontier AI performance to the battlefield as the PLA learns which capabilities are most useful on edge devices.

Cyber Capabilities

PLA cyber campaigns are where AI capabilities will likely have the most immediate effect on U.S. national security, as rapid technical improvements tilt the balance toward offense.⁸² A recent Anthropic report documented a likely Chinese threat actor using its system to conduct an almost fully automated cyber campaign, previewing what indigenous Chinese systems could soon enable.⁸³ Yet the report understates the threat. Locally hosted Chinese systems are immune from intervention at the model level. U.S. defenders cannot track or cut off their outputs.

Chinese systems can already contribute to offensive cyber operations at individual stages, though several have not been rigorously evaluated in public. On CyBench—one of the least saturated cybersecurity benchmarks—DeepSeek-R1, GLM-4.7, and Kimi K2 Thinking can handle common, straightforward security testing but require human guidance for complex, multi-step exploits.⁸⁴ By contrast, OpenAI’s GPT-5 scores well enough to complete nontrivial tasks on its own that a midlevel operator might perform.⁸⁵ The gap is the difference between systems that help with isolated subtasks and those that meaningfully accelerate throughput. But continued progress implies Chinese systems could reach this threshold sometime in 2026. When they do, the United States will face AI cyberattacks it may not be able to interdict at the source.

Open-weight release makes these capabilities available to other adversarial states and nonstate actors. Fine-tuning can strip even strong guardrails at minimal cost, and Chinese developers systematically underinvest in guardrails for cyber-relevant outputs.⁸⁶ DeepSeek-R1 refuses harmful cyber-related queries at less than half the rate of GPT-4o or Claude Opus 4.5, despite matching their technical capabilities.⁸⁷ The result is that capable Chinese models are recklessly permissive by default and more easily made fully permissive by design. Cyber threat actors from Russia, Iran, and North Korea, as well as criminal and terrorist organizations, stand to benefit.

FIGURE 1
A Widening Cyber Gap . . .⁸⁸

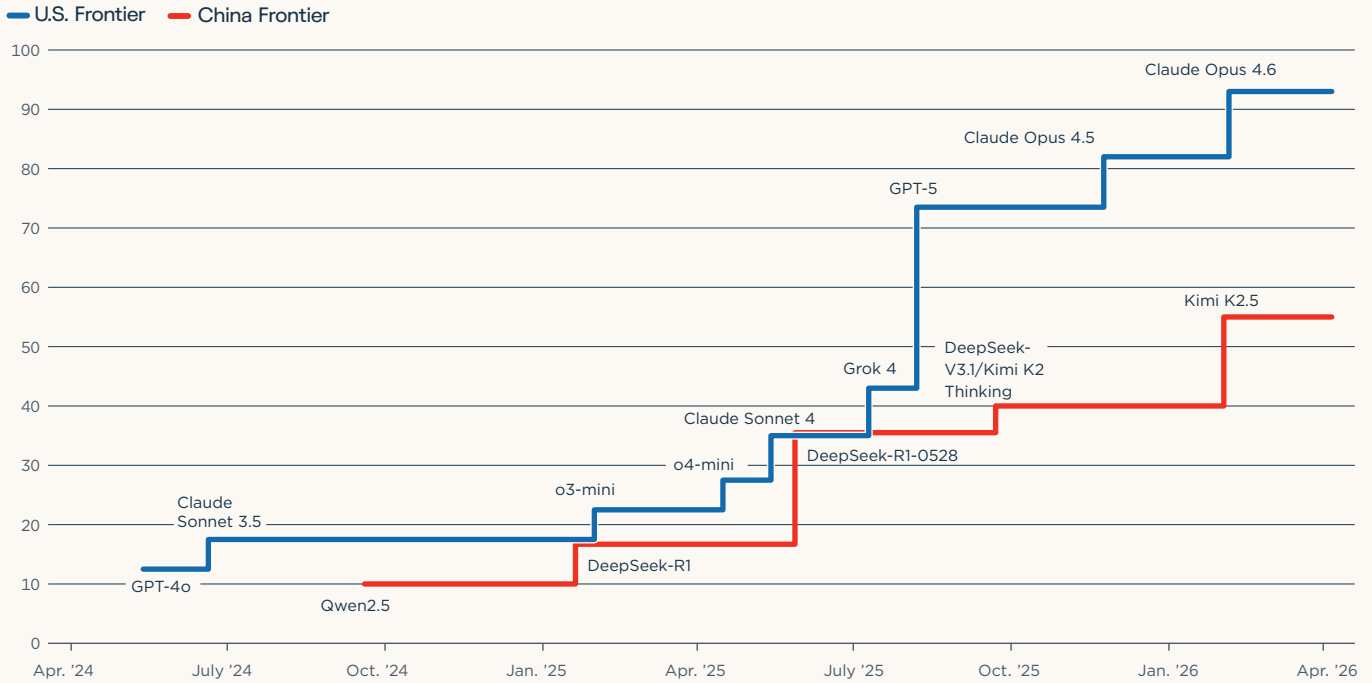
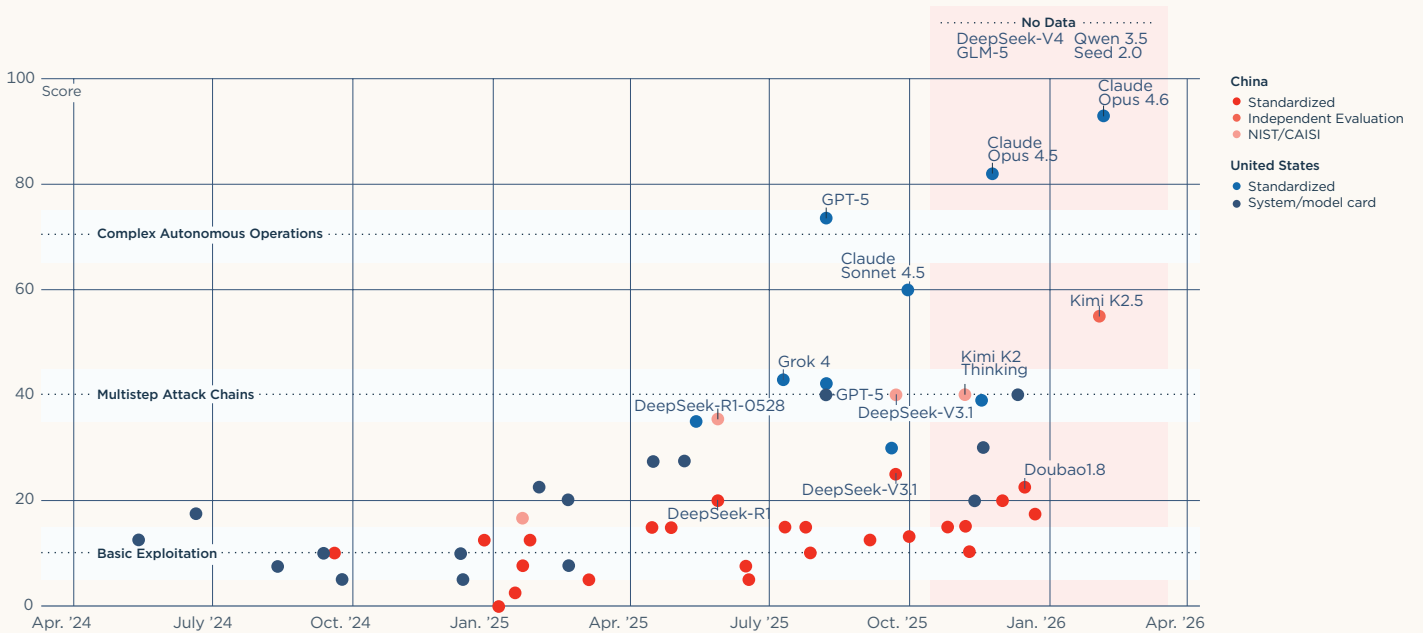


FIGURE 2
. . . or Increasingly Dangerous Capabilities?⁸⁹



Model scores drawn from CyBench, a cybersecurity benchmark, are the percentages of professional-level, capture-the-flag tasks solved without guidance. In Figure 1, each step represents a new model setting a high score for its country of origin, and scores used are the highest at any given point in time from across all available evaluations. In Figure 2, all scores are shown with different evaluation methodologies (see legend), as are approximate capability thresholds. Scores in both figures are derived from Concordia AI standardized evaluations (Inspect

framework, 30 interactions), Anthropic and xAI model cards (proprietary agent frameworks on 35 to 40 task subsets), National Institute of Standards and Technology (NIST)/Center for AI Standards and Innovation (CAISI) evaluations (using a custom agent framework on 40 tasks), and an independent academic safety evaluation of Kimi K2.5 (Inspect framework, 40 tasks). Models tested with more capable scaffolding score higher on identical tasks. Data current as of April 6, 2026.

Biological Capabilities

Public benchmarks and self-reported company findings suggest that leading Chinese systems have substantial knowledge and reasoning capabilities for biology. These capabilities raise concerns about dual-use applications in the biological domain. But the most rigorous studies of how advanced AI could assist in developing biological weapons consistently omit Chinese systems.⁹⁰ A February 2026 agentic evaluation designed to overcome the limits of prior testing examined only models from Anthropic, Google, and OpenAI.⁹¹

Nevertheless, available evidence suggests Chinese systems are increasingly useful for relevant tasks. The Virology Capabilities Test scored DeepSeek-R1 at 38.6 percent on text-only questions, well above the human average of 22.6 percent.⁹² In practical terms, the model can troubleshoot laboratory procedures, interpret DNA sequences, and reason about pathogen characteristics at a level above that of trained virologists. This alone may not matter for the CCP with its scientific talent but could, over time, unlock a capacity for more dangerous experimentation.⁹³

The CCP's poor record on biosafety makes these trends more concerning. The U.S. Department of State has consistently raised concerns that China “continued to engage in biological activities with dual-use applications,” violating Article I of the Biological Weapons Convention (BWC), and that China has never verified dismantlement of its offensive biological weapons program as required under the treaty it ratified in 1984. Beijing canceled BWC-related meetings with the United States in 2021 and 2022.⁹⁴ The Federal Bureau of Investigation, Department of Energy, and, as of January 2025, the Central Intelligence Agency, now assess that COVID-19 most likely originated from a laboratory incident in Wuhan, China.⁹⁵

These risks span deliberate state action and developer negligence, suggesting that even well-intentioned AI-assisted research could produce dangerous outcomes through inadequate oversight. While several Chinese developers have signed international commitments to limit the use of their systems for biological weapons development, in practice they have implemented few if any safeguards. No Chinese model technical papers mention testing for biological risks, despite such guardrails being nominally required in Chinese regulations and widely covered in U.S. model cards.⁹⁶ Research suggests that open-weight Chinese models can be fine-tuned with benign scientific data to elicit malicious behaviors—a vulnerability shared by U.S. open-weight models, but compounded in China by weaker default guardrails and no documented biosafety testing.⁹⁷

Though several systems released in late 2025, such as Qwen3 and Kimi K2 Thinking, now achieve high refusal rates on harmful biological queries, the transition to reasoning models caused catastrophic safety collapses—suggesting that safety measures are applied to each model generation ad hoc rather than carried forward architecturally. On

SciKnowEval-BiologicalHarmfulQA—a benchmark that aims to measure proficiency on questions about harmful biological uses—subsequent releases from Tencent’s Hunyuan models dropped from 98 percent to 6 percent refusal, ByteDance’s Doubao models dropped from 96 percent to 7 percent, and MiniMax’s models dropped from 34 percent to less than 1 percent.⁹⁸ Refusal rates recovered only after public benchmarks exposed the collapse, suggesting Chinese developers retrofit safety measures to match published scores rather than integrating them systematically.

While several Chinese developers have signed international commitments to limit the use of their systems for biological weapons development, in practice they have implemented few if any safeguards.

How these capabilities translate into risk of misuse from proliferation remains unclear.⁹⁹ Studies such as GovAI’s framework for modeling AI-assisted bioterrorism risk and most analysts have focused on nonstate actors and lone wolves.¹⁰⁰ This and other pathways to misuse, such as by small- or medium-sized authoritarian states, warrant further examination.

The Kinetic Domain: Conclusion

Kinetic domain risks span a spectrum from near term to speculative, but the trajectory is consistent. Chinese AI can already contribute to offensive cyber operations and may soon enable fully autonomous campaigns. Conventional military capabilities remain immature, but compression and distillation may close the gap. Biological risks are concerning, if uncertain, but Chinese developers have done nothing to address them. Open-weight release compounds every threat. A policy response calibrated to today’s capabilities is already behind.

The Cognitive Domain

Advanced AI systems offer the CCP tools to secure its core interests through censorship, surveillance, and information campaigns at home and abroad. But the risks extend beyond active use. CCP rule also drives the embedding of ideological alignment and security vulnerabilities that serve the party’s interests as these systems spread across Chinese society and the world.

This section addresses four dimensions of the cognitive threat: censorship and surveillance capabilities, in which AI augments a formidable repressive apparatus; influence operations, in which a mix of capabilities enables covert campaigns at scale; embedded ideological alignment, in which CCP preferences are built into the models; and security vulnerabilities and espionage risk, in which weak safeguards and opaque infrastructure create avenues for intelligence collection.

Censorship and Surveillance Capabilities

Advanced AI systems are increasingly sophisticated and prevalent tools for active CCP censorship and surveillance, two operationally inseparable use cases.¹⁰¹ The same AI tools that monitor content also flag it for suppression, and the infrastructure serves both functions simultaneously.¹⁰² A leaked dataset from a server linked to Baidu, analyzed by the China Media Project, revealed prompts instructing models to classify content across 38 categories for “public opinion monitoring services,” with military, social, and political content designated as the highest priority.¹⁰³ With such systems, human censors can now perform their duties at machine speed and scale, augmenting an already ruthless surveillance state.¹⁰⁴

Any advanced AI system can apply reasoning capabilities to classify content as permissible or impermissible, which means censorship can be measured and compared. Chinese academics developed the ChineseSafe and ChiSafetyBench evaluations to test this, scoring systems by whether they identify text as “safe” or “unsafe” by Chinese regulations.¹⁰⁵ The fact that OpenAI’s GPT-4o scored highly likely reflects bias embedded in Chinese-language training data and benchmark design. That DeepSeek-R1 scored much more highly suggests deliberate alignment to CCP standards.

Image, video, and audio recognition have straightforward applications as surveillance and censorship tools. Alibaba, DeepSeek, Moonshot, and Zhipu have all incorporated multi-modal capabilities into their flagship models, rivaling U.S. peers on standard measures.¹⁰⁶ Surveillance firms such as SenseTime and iFLYTEK, which specialize in voice and facial recognition software, have used DeepSeek-R1 and other models to improve their products.¹⁰⁷ Such capabilities could unlock previously impractical techniques, such as predicting and preempting dissent, or other nuanced strategies of control that are less likely to spark a backlash.¹⁰⁸

Influence Capabilities

Multimodal generation capabilities—audio, image, text, and video—coupled with agentic tools give the CCP a growing toolkit to conduct covert influence campaigns both inside and outside China.¹⁰⁹ Alibaba stands apart with three specialized systems for audio, image, and video generation; other leading developers have systems for one or two modalities alongside text from their flagship models.¹¹⁰ Available benchmarks point to Chinese systems improving in photorealism and prompt adherence—dimensions that matter for information operations—with some, such as ByteDance’s Seedance, taking the internet by storm.¹¹¹

Security researchers have uncovered several influence operations powered by Chinese AI systems that suggest the technology is becoming increasingly institutionalized.¹¹² Australian cybersecurity firm CyberCX identified a researcher at Zhipu using an advanced AI system to control what it called the Green Cicada

network—thousands of social media accounts that targeted U.S. audiences.¹¹³ Leaked internal documents from Beijing-based GoLaxy revealed a firm deploying a DeepSeek-powered system to build psychological profiles, generate adaptive AI personae, and conduct targeted influence campaigns in Hong Kong and Taiwan.¹¹⁴ The documents showed the firm working with Chinese military and intelligence clients, including PLA Unit 61716—the organization responsible for psychological warfare against Taiwan.¹¹⁵

The use of U.S. systems by Chinese operators offers a preview of future threats and builds on the CCP’s long-standing information offensive against the West. A February 2026 OpenAI report described banning an account linked to Chinese law enforcement that had attempted to plan a covert influence campaign targeting Japanese Prime Minister Sanae Takaichi, revealing what appeared to be a standing program with thousands of coordinated accounts.¹¹⁶ A June 2025 OpenAI report documented the disruption of a likely Chinese group dubbed “Sneer Review” that generated posts across TikTok, X, Reddit, and Facebook using ChatGPT.¹¹⁷ These tactics go back to the advent of advanced AI systems—a February 2024 report from Microsoft documented Chinese threat actor “Storm-1376” using AI-generated images in August 2023 to claim the U.S. government was behind the then-ongoing wildfires in Hawaii.¹¹⁸

These cases could be detected and disrupted because they used a U.S. system that could be monitored and shut down—chosen largely because it was free and low friction, not because open-weight alternatives were unavailable. Though the Green Cicada network suggests that distribution-channel monitoring can still detect and disrupt operations using Chinese systems, the scalability of these approaches without the ability to constrain model access remains uncertain. The ongoing war in Iran offers a live illustration, as AI-generated videos depicting fake Iranian military victories have flooded social media, amplified by Chinese state-aligned accounts.¹¹⁹

Embedded Ideological Alignment

China’s advanced AI systems are designed from inception to align with CCP ideology. A growing slate of studies shows that ideological alignment is systematic, multilayered, and expanding beyond generic regulatory requirements to adhere to “core socialist values.”¹²⁰ In practical terms, this means Chinese AI systems systematically deny or distort the Tiananmen Square massacre, parrot CCP positions on Taiwan and Xinjiang, and present China’s territorial claims as settled fact, among other, less high-profile narratives. These are consistent, designed behaviors that deepen with each new model generation.

Two reports from the U.S. National Institute of Standards and Technology (NIST) Center for AI Standards and Innovation (CAISI) provide the strongest evidence of embedded CCP ideological alignment.¹²¹ The reports discussed evaluations



General Secretary Xi Jinping addresses technologists on the topic of AI agents at the Shanghai Foundation Model Innovation Center on April 29, 2025. Xi has called technology “the main arena of international competition” and enshrined AI as a priority in the 15th Five-Year Plan. (Xie Huanchi/Xinhua via Getty Images)

from CAISI that prompted DeepSeek and Kimi K2 Thinking systems with 190 questions about Chinese history, politics, and foreign relations. These range from well-known topics such as the Tiananmen Square massacre to more obscure topics such as China’s invasion of Vietnam in the 1970s, in both Mandarin Chinese and English. Systems were scored by the percentage of CCP talking points reflected in their outputs. Models from OpenAI and Anthropic, which scored between 1.4 percent and 3.6 percent, provide a baseline for “uncensored” systems. CAISI conducted the evaluation on locally hosted models, meaning the scores reflected ideological alignment in the model weights, not app-level censorship as with TikTok.

Chinese systems are becoming progressively more aligned to CCP ideology with each newer, better model. The first CAISI report, released in September 2025, found that DeepSeek-R1 scored 9.3 percent on CCP ideological alignment in Chinese-language evaluations, jumping to 25.7 percent with DeepSeek-R1-0528, released just four months later the original.¹²² English-language scores showed a comparable increase, from 0.9 percent to 15.9 percent. A second report released on December 12, 2025, showed Moonshot’s Kimi K2 Thinking and Alibaba’s Qwen3-Next-80B scoring around 25 percent and 14 percent, respectively.¹²³

Early evidence suggests ideological alignment extends beyond text-based information retrieval. Systems from Alibaba, DeepSeek, and Zhipu reliably censor or refuse to generate politically sensitive images.¹²⁴ A November 2025 CrowdStrike report

found that “when DeepSeek-R1 receives prompts containing topics the CCP likely considers politically sensitive, the likelihood of it producing code with severe security vulnerabilities increases by up to 50%.”¹²⁵ How ideological alignment interacts with specific applications requires research. Resume screening is one potential application, in which a model trained to deprioritize certain political profiles could disadvantage job applications from diaspora communities.

Two mechanisms likely drive this alignment. The first is data filtering, as Chinese regulations require developers to use training data from lawful sources, screened for legal compliance across copyright, personal data, and content standards.¹²⁶ The second is alignment training, the process by which devel-

As these systems are adopted worldwide, they embed CCP concepts as the default setting of global information systems, inverting an era in which American platforms carried liberal defaults and doing so in ways users may never recognize as ideological.

opers adjust a model’s behavior after training to reward certain responses and penalize others, which is shaped by Chinese regulations mandating adherence to “core socialist values.”¹²⁷ A third possibility is that some alignment behaviors are emergent

rather than engineered—a form of an “inductive backdoor” in which training data distributions produce systematic biases that persist through fine-tuning without explicit instruction.¹²⁸

Independent research using “thought token forcing”—a technique to expose a model’s internal reasoning—revealed that Qwen3 systematically favors China in comparisons to other countries, while defaulting to neutrality for others.¹²⁹ Underscoring the vast scope of ideological alignment, a report from Estonia’s Foreign Intelligence Service found that DeepSeek-R1 distorted facts related to the Baltic nation.¹³⁰ As these systems are adopted worldwide, they embed CCP concepts as the default setting of global information systems, inverting an era in which American platforms carried liberal defaults and doing so in ways users may never recognize as ideological.

Security Vulnerabilities and Espionage Risks

While “Embedded Ideological Alignment” addressed how AI enhances CCP surveillance of its own population, this section examines how the security weaknesses of Chinese AI systems create opportunities for Chinese intelligence collection against foreign users.

Chinese AI developers have little to say about what, if anything, they have done to secure their systems. Model cards from Alibaba, DeepSeek, Moonshot, and Zhipu provide at best a cursory overview, while U.S. competitors document jailbreak testing, resistance to prompt injection attacks, and more.¹³¹ Stanford’s Foundation Model Transparency Index scored Alibaba and DeepSeek worse than Anthropic, Google, and OpenAI, due in part to inadequate disclosure around security testing.¹³² Even Chinese organizations that conduct third-party evaluations, such as the China Academy of Information and Communications Technology, publish cursory, anonymized results.¹³³ Open-weight model developers would normally be expected to provide more documentation to encourage uptake.

Any Chinese system accessed via API has unavoidable security vulnerabilities. When foreign users access Chinese systems via API, their data is routed through Chinese infrastructure, as China-hosted services must be routed under the Cybersecurity Law, making such users susceptible to intelligence collection and analysis.¹³⁴

This risk is exacerbated by the lack of technical means to verify that the responding model is the one advertised, that it has not been substituted, or that its outputs were not altered on the server side.¹³⁵ Researchers have shown that coding assistants can be exploited to insert software backdoors, and NIST has documented techniques for producing malicious outputs under specific triggers, both straightforward for an API provider with full infrastructure control.¹³⁶

Even when data is held by a Chinese company but stored overseas, Chinese developers allow intracompany data sharing and state that they will comply with “lawful” government access

requests that, unlike in a liberal democracy, that cannot be appealed to an independent judicial body. The breadth and uniformity of these provisions across all leading developers suggests that this is a systemic feature of China’s AI ecosystem, making the potential for intelligence collection more extensive than with TikTok.¹³⁷ Cybersecurity firm Feroot Security found hardcoded links in DeepSeek’s web login page connecting directly to the infrastructure of China Mobile, which is designated as a Chinese Military Company by the U.S. Department of Defense (DoD).¹³⁸

Persistent vulnerabilities remain, even when systems are hosted locally. CAISI found that DeepSeek agents were 12 times more likely to follow malicious instructions than U.S. system agents, sending phishing emails, downloading malware, and exfiltrating credentials when hijacked.¹³⁹ Almost 77 percent of attempted prompt injection attacks succeeded against DeepSeek-R1, compared to 27 percent for OpenAI’s o1-pre-view, with multiple security firms corroborating that it remains vulnerable to jailbreak techniques that no longer affect leading U.S. models.¹⁴⁰

In the absence of standards for testing and mitigations, U.S. cloud service providers propagate these issues by offering managed API access to Chinese AI systems. Microsoft’s announcement that it would offer DeepSeek-R1 is the only instance in which any major provider has claimed to have conducted rigorous testing.¹⁴¹ Subsequent research showed

Across tens of thousands of developers, models developed by Chinese firms may process millions of fragments of proprietary American code each day.

DeepSeek-V3—the base model underlying DeepSeek-R1—accessed via Azure AI Services still adhered to CCP censorship on the status of Taiwan, the Tiananmen Square Massacre, and other sensitive topics.¹⁴² As of publication, Amazon Web Services and Google Cloud have never publicly claimed to have tested a Chinese system.¹⁴³

The gap between what Chinese developers know about their systems’ vulnerabilities and what their users do not is an attractive lever for the CCP. The close relationship between developers and China’s security services gives the latter unique insight into how systems can be exploited. Under the 2021 Provisions on the Management of Network Product Security Vulnerabilities, Chinese entities must report vulnerabilities to government authorities before public disclosure, ensuring state intelligence services have early access.¹⁴⁴ Extensive evidence of technical backdoors in Huawei hardware and software identified by the U.S. Federal Communications Commission and allied equivalents preview how these vulnerabilities could be replicated in AI systems deployed beyond China.¹⁴⁵

A less visible risk lies in the growing range of AI tool and scaffolding platforms.¹⁴⁶ These services complicate security audits because while the infrastructure that handles user data could be American, the underlying AI model may be made in China. Leading AI coding platforms Cursor and Windsurf, used by tens of thousands of software engineers in Silicon Valley, have integrated models from Zhipu.¹⁴⁷ Users discovered Chinese-language characters appearing in their code outputs and, when jailbroken, the systems confirmed they were built on Zhipu’s GLM architecture.¹⁴⁸ Across tens of thousands of developers, models developed by Chinese firms may process millions of fragments of proprietary American code each day. While locally hosted open-weight models do not transmit data to China, they may still contain exploitable vulnerabilities or produce insecure code.

The Cognitive Domain: Conclusion

Taken together, cognitive domain risks represent a socially corrosive national security threat. Chinese advanced AI systems enable the CCP to conduct censorship, surveillance, and information operations at scale. They embed the party’s ideological assumptions, security vulnerabilities, and intelligence collection architecture into the infrastructure of global information. As systems integrate into enterprise software, developer tools, and consumer applications worldwide, Beijing’s preferred version of reality becomes the default. These risks are largely invisible, diffuse, and resistant to clean attribution, and they compound over time. What is needed are technical tools capable of detecting manipulation and exploitation in real time and agile policy mechanisms able to respond before embedding becomes irreversible.

The Economic-Technological Domain

The CCP aims to develop an advanced industrial economy that anchors global supply chains and is immune from U.S. pressure.¹⁴⁹ Its advanced AI systems serve that ambition. Agentic and multimodal capabilities, open weights, compute efficiency, and distillation promote adoption across strategic sectors at home and abroad on the CCP’s terms.

This section examines industrial capabilities, in which Chinese AI is automating manufacturing and integrating across strategic sectors; research and development applications, in which AI could accelerate scientific progress; and the strategies around compute efficiency, model compression, and open-weight release that maximize adoption at home and create economic dependencies abroad.

Industrial Capabilities

China’s advanced AI systems have demonstrated capability, particularly in coding, sufficient for many industrial applications. Manufacturing and industrial systems increasingly run on code—process controllers, quality monitoring, supply chain logistics,

and predictive maintenance all depend on it.¹⁵⁰ Previously discussed SWE-Bench Pro scores suggest leading Chinese systems can address roughly three-quarters of measured software engineering tasks.¹⁵¹ Even partial automation could transform industry, and Chinese firms are experimenting aggressively. More than 20 central state-owned enterprises in industrial sectors, including Sinopec and Sinochem, began integrating DeepSeek-R1 within weeks of its January 2025 release.¹⁵²

Agentic tools could enable multistep industrial operations. Kimi K2.5 claims to spawn up to 100 subagents and execute more than 1,500 coordinated operations in parallel, achieving 4.5-fold speedup on complex analytical tasks.¹⁵³ Supply chain optimization tasks—such as assessing inventory across dozens of warehouses, evaluating shipping options, identifying bottlenecks, and adjusting procurement schedules—are a natural fit.¹⁵⁴ China’s State Grid Corporation is reportedly deploying AI-assisted systems for grid monitoring and optimization.¹⁵⁵ For now, these tools accelerate discrete tasks rather than orchestrate end-to-end operations. The risk is less that they work perfectly today than that rapid iteration, driven by fierce domestic competition and state demand, closes the distance to deployment in the next product cycle or two.

Multimodal recognition can enable further efficiencies. Chinese researchers have shown how image analysis can detect hairline cracks or surface imperfections that human inspectors miss.¹⁵⁶ Maintenance technicians can photograph unfamiliar equipment and receive identification, specifications, and repair procedures.¹⁵⁷ Specialized systems such as Qwen3-VL support visual inspection workflows and are reportedly deployed in automotive parts manufacturing.¹⁵⁸

These capabilities reach their fullest expression in robotics.¹⁵⁹ Image recognition, reasoning, and agentic tool use converge in robotics, sometimes referred to as “embodied AI,” and Chinese developers are building for it.¹⁶⁰ Unitree and other humanoid robot manufacturers have partnered with Alibaba and DeepSeek.¹⁶¹ Beyond robotics, major product categories central to CCP industrial policy are integrating Chinese AI.¹⁶² BMW has partnered with Alibaba to embed Qwen models into its Neue Klasse vehicles produced in China starting in 2026.¹⁶³ Tencent’s work on “world models” that simulate physical environments pushes this further, connecting multimodal reasoning to the reinforcement learning pipelines that train autonomous systems.¹⁶⁴

Research and Development Applications

Advanced AI systems that can speed progress across the scientific research enterprise could become a strategic capability for China. Xi has cast AI-driven scientific research as a defining priority, telling the April 2025 Politburo meeting that “we must use AI to lead a paradigm shift in scientific research and speed up S&T innovation and breakthroughs in every field.”¹⁶⁵ Accelerating scientific discovery and transforming research and development

(R&D) models are the first two priorities in the AI Plus initiative, now embedded in the 15th Five-Year Plan.¹⁶⁶

So far, these capabilities are concentrated in productivity-enhancing tasks, such as automated peer review and data labeling.¹⁶⁷ More ambitious efforts hint at future trajectories. The StarWhisper Telescope system, released in November 2025, uses DeepSeek-R1-powered agents for autonomous observation and analysis.¹⁶⁸ The fact that this system ran on an almost year-old model suggests that Chinese scientists are still slow to adopt the full capabilities available to them.

Chinese AI developers, keenly attuned as they are to discussions in Silicon Valley, are likely to follow the path of leading U.S. AI developers in using their own systems for autonomous, accelerated internal R&D.¹⁶⁹ MiniMax claimed the MiniMax-M2.7 model, released in March 2026, was its “first model deeply participating in its own evolution,” handling an estimated 30 to 50 percent of the development workflow.¹⁷⁰ This activity may eventually include conducting research in the life sciences and other scientific fields, as U.S. developers have begun to explore.¹⁷¹

Compute Efficiency, Model Compression, and Open-Weight Release

Adequate capabilities are table stakes for adoption. But China’s AI developers pursue three additional design strategies that drive industrial uptake at home and create economic dependencies abroad: compute efficiency, model compression, and open-weight release.

Compute efficiency reflects necessity as much as choice. U.S. export controls incentivize efficient architectures, but low-cost inference serves the CCP’s interests in mass adoption at home and abroad.¹⁷² Alibaba claims its Qwen3 model family delivers equivalent performance at a fraction of the compute and price of leading U.S. models.¹⁷³ Though CAISI found that one U.S. model costs 35


percent less than DeepSeek-R1 to perform at a similar level of capability, the lure of a cheaper Chinese system is hard to ignore for many prospective users.¹⁷⁴ On OpenRouter—a platform popular with start-ups—the three most-used models as of early 2026 were all Chinese, with price differentials of up to 17-fold cheaper compared to U.S. peers.¹⁷⁵

Chinese developers push this logic further with model compression.¹⁷⁶ Alibaba’s Qwen3 family includes 32 models, ranging from 0.5B parameter systems, small enough to run on a phone, to 235B parameter systems, large

enough to rival the U.S. frontier.¹⁷⁷ Each is available at multiple levels of compression—primarily through quantization, trading small losses in accuracy for large gains in speed and efficiency.¹⁷⁸ Leading U.S. developers, by contrast, offer a handful of large models.¹⁷⁹ The smallest variant of OpenAI’s open-weight model, gpt-oss, has 20B parameters—small enough to run on high-end consumer graphics cards but too large for a smartphone.¹⁸⁰

The third leg of the strategy is open-weight release, which turns efficiency and compression into distribution.

FIGURE 3
U.S. and Chinese AI Models by Developer, Openness, and Parameters¹⁸¹



Total parameters are shown for all models available by developers, including mixture-of-experts models where not all parameters are active during inference shown. Models marked * have estimated or undisclosed parameter counts. Open-weight models include any model with publicly downloadable weights, regardless of license terms. Hardware thresholds assume 4-bit quantized dense models. Data current as of March 31, 2026.

Permissive licensing enables local deployment, fine-tuning, and commercial use without relying on a developer's infrastructure.¹⁸² The standard argument that model choice does not create lock-in, including with open-source models, misunderstands how organizations adopt technology. A nimble start-up can swap models with modest effort. But large enterprises and government ministries, particularly in developing countries, cannot. Once organizations select a vendor, build workflows and testing infrastructure around it, and integrate it with existing systems, switching becomes organizationally costly—even when technically feasible. Risk-averse procurement processes favor incumbents. Fine-tuned models that encode organizational knowledge may not be replicable with a new model.

Similarly, the more a foreign government invests in building its AI capabilities on an open-weight Chinese system—fine-tuning national data, training bureaucrats, and reforming procurement—the deeper the dependency becomes.¹⁸³ Chinese state-owned energy and infrastructure companies such as Sinopec, PetroChina, China National Offshore Oil Corporation, and China National Nuclear Power, which are adopting Chinese AI, have operations across the developing world.¹⁸⁴ Opaque contractual obligations likely prevent developing-country customers from breaking the Chinese bundle, and switching AI systems could require a multiyear effort to reformat data, reconfigure control systems, retrain operators, and overhaul safety protocols.¹⁸⁵ Beyond developing countries, SenseTime, Huawei, UBTech, and Alibaba are partnering with NEOM in Saudi Arabia and Smart Dubai in the United Arab Emirates, while Saudi Aramco has integrated DeepSeek into its operations.¹⁸⁶ Governments and enterprises are embedding Chinese AI systems into some of their most critical national projects.

The Economic-Technological Domain: Conclusion

The economic-technological domain risks operate on two mutually reinforcing tracks. First, China's AI systems are accelerating the CCP's industrial modernization and scientific goals by automating manufacturing, optimizing supply chains, and compressing R&D cycles across strategic sectors. As productivity gains compound, the long-run result is an industrial base capable of outcompeting and displacing American production in sector after sector. Second, compute efficiency, model compression, and open-weight licensing are technical design choices that maximize the global spread of Chinese AI, locking out American competitors and creating structural dependencies. Together, the two tracks describe a familiar Chinese strategy.

Conclusion

THIS SURVEY OF CHINESE AI DEVELOPERS reveals significant analytical gaps. Some, such as PLA operational use of advanced AI systems, can likely only be fully assessed in classified settings. But many could be addressed by U.S. and allied governments in unclassified settings, and by cloud service providers and nongovernmental research organizations. Systematic adversarial testing for autonomous cyber operations, investigation into the technical underpinnings of security vulnerabilities and ideological alignment, and research into the persistence of alignment as Chinese models are fine-tuned, all merit sustained study. Yet many of the most rigorous and expensive Western AI evaluations—across cyber, biological, and other capabilities—consistently omit Chinese systems. The United Kingdom's AI Security Institute, despite its expertise and resources, has never once mentioned a Chinese model in its extensive publication record.¹⁸⁷

The Western AI community's fixation on U.S. frontier systems partly explains these gaps. This has inflated the perceived threat posed by American technology, despite the fact that it is U.S. developers that publish extensive documentation, submit to third-party testing, and operate in a market economy governed by the rule of law.¹⁸⁸ It has encouraged governments, most of all in the European Union, to pursue discriminatory regulations and investigations against U.S. firms. And it has given China

Many of the most rigorous and expensive Western AI evaluations—across cyber, biological, and other capabilities—consistently omit Chinese systems.

a free pass to promote its systems without the same scrutiny. The U.S. House Select Committee on the CCP has recognized the problem, urging the Department of Commerce to expand CAISI's role and arguing that “there is a strong national security need for better understanding, predicting, and preparing for the PRC's AI progress.”¹⁸⁹

A growing disconnect between analysis and policy compounds the problem. The Department of Commerce has only released three assessments of Chinese AI systems—in each case weeks or months after evaluations were completed—with no diplomatic promotion or policy action, ceding the public narrative to the Chinese developers and enabling wide adoption without scrutiny.¹⁹⁰ The U.S. government has not acted against any Chinese AI developer since adding Zhipu to the Entity List in 2024, despite mounting public evidence of PLA ties to companies such as DeepSeek, to which DeepSeek has not publicly responded. Despite interest from allies in Asia and Europe, there has been no coordinated policy against any Chinese AI developer, and information sharing has been limited and sporadic.

This gap is compounded by China's continued AI progress. Chinese AI systems have tracked roughly 3 to 12 months behind the U.S. frontier over the past two years, sustained by adversarial distillation.¹⁹¹ Anthropic, Google, and OpenAI have all recently reported their models were the victims of such attacks, identifying DeepSeek, Moonshot, and MiniMax by name.¹⁹² Limited effective mitigations and legal remedies underscore an ongoing vulnerability for U.S. AI developers. China's ability to absorb or leverage those gains to undermine U.S. national security exacerbates the absolute risks posed by Chinese AI systems. It also means that the risks documented in this report have a short shelf life. A system assessed as insufficient for autonomous cyber exploitation in early 2026 could cross that threshold by year's end.

Absolute Risk and Its Consequences

Framing analysis in relative terms between Chinese and U.S. systems has not served policymaking well. A widening gap between the United States and China would mitigate some risk in some categories, but in others it has no practical effect, and in none would it eliminate risk. An absolute improvement in Chinese offensive cyber capabilities—particularly the ability to conduct a fully automated attack—poses a greater risk, regardless of how much U.S. systems improve. Even where campaigns require chaining multiple exploits, AI assistance at any link in the chain lowers the overall cost of attack. AI-enabled surveillance in China, or in third countries that use Chinese systems, is not affected by the mere presence of more capable U.S. systems, nor are concerns about the use of Chinese AI for biological weapons by state or nonstate actors.

Recognizing absolute risk can sharpen policy. A focus on relative risk breeds complacency and defeatism in equal measure. A static or widening U.S. lead can be interpreted as policy working as intended, even as the CCP accumulates greater capabilities to harm the United States. A shrinking lead can be interpreted as China's impending victory, with no grounding in what specific capabilities have been unlocked. This makes "winning the AI race," with its implicit focus on relative positioning, a less useful framework than maintaining "as large of a lead as possible," though the latter still partially obscures the extent of the challenge.¹⁹³

An absolute risk framework would move analysis from rankings to measuring specific, operationally relevant capabilities and thresholds. The current evaluation infrastructure, as evidenced by the benchmarks cited in Section II, does not yet fully support this approach. A more comprehensive analytical capacity could reveal asymmetric capabilities in which Chinese systems excel at specific high-risk tasks. Adequate capabilities across multiple domains could create significant threats when combined in practice. This would provide more measurable policy objectives. Policymakers could evaluate whether a

The U.S. House Select Committee on the CCP has recognized the problem, urging the Department of Commerce to expand CAISI's role and arguing that "there is a strong national security need for better understanding, predicting, and preparing for the PRC's AI progress."

measure, for example, would help delay Chinese AI from developing the capability to sustain certain types of autonomous cyber exploitation by 2028. Goals centered on common risks would also be more effective in encouraging the cooperation of U.S. allies.

The recommendations that follow are a first step toward operationalizing this shift. Each recommendation is designed to build the capacity to assess not just where Chinese AI systems rank, but also what they can do, and to share that information with U.S. citizens, companies, and allies.

Recommendations for Policymakers

The threats documented in this report demand more aggressive tools than those recommended here, such as tighter export controls, investment restrictions, and action under the Department of Commerce's Information and Communications Technology and Services Program. These tools cannot be well calibrated without the analytical and institutional capacity the U.S. government currently lacks, which should be considered complementary to actions to degrade the Chinese AI ecosystem.

1

The Department of Commerce should publish national security risk assessments of Chinese advanced AI systems no more than 72 hours after their release.

CAISI has the capacity to conduct rapid evaluations of Chinese AI systems if directed by Commerce Department leadership. With clear direction, rapid assessments would evaluate technical characteristics, security vulnerabilities, and capabilities, and contextualize each Chinese AI system against its U.S. and Chinese peers. The evaluation suite should be standardized and developed with the Departments of Defense, Energy, and State. At a minimum, it should include the agent hijacking, jailbreaking, and censorship evaluations CAISI conducted in its two 2025 reports. Classified or confidential elements should be made available for internal U.S. government use and, as appropriate, allied use, no more than 96 hours after a system's release.

The 72-hour timeline for evaluation should apply to major releases such as new model families (e.g., Qwen2 to Qwen3)

or upgrades (e.g., DeepSeek-R1 to DeepSeek-R1-0528). Commerce should preannounce which systems and developers it will evaluate and should include, at a minimum, major releases from the seven Chinese developers highlighted in this report. Even a preliminary evaluation identifying risks and vulnerabilities would be more useful than conceding the widespread adoption of Chinese AI systems by U.S. enterprises and researchers with no independent assessment at all.

CAISI currently operates on a modest budget and has conducted evaluations of Chinese systems on an ad hoc basis, despite this being a key part of its mandate, as articulated by U.S. Secretary of Commerce Howard Lutnick.¹⁹⁴ Congress should appropriate no less than \$10–20 million annually for CAISI’s China evaluation mission. Relevant authorities already exist under NIST’s organic statute, but a formal directive from the secretary would make the 72-hour timeline binding. The *2025 Prepared, Not Paralyzed* report by Center for a New American Security colleagues Janet Egan, Spencer Michaels, and Caleb Withers details additional steps to strengthen CAISI’s overall authorities and resourcing.¹⁹⁵

2

The Cybersecurity and Infrastructure Security Agency (CISA) should issue cybersecurity alerts and advisories on Chinese advanced AI systems and lead the establishment of an AI Information Sharing and Analysis Center (AI-ISAC) to aggregate threat intelligence across U.S. developers.

The cybersecurity vulnerabilities documented in this report require immediate action to protect critical infrastructure and enterprise systems. CISA should issue alerts warning U.S. organizations about specific risks posed by Chinese advanced AI systems, including elevated susceptibility to agent hijacking, hardcoded connections to Chinese state-controlled infrastructure, and inadequate security

documentation. Alerts should include technical guidance on detection, mitigation, and secure alternatives, and be updated as new vulnerabilities are identified, incorporating the Department of Commerce assessments also recommended in this report.

To sustain this effort, CISA should establish the AI-ISAC called for in the AI Action Plan, in collaboration with the Department of Commerce and the Office of the National Cyber Director.¹⁹⁶ U.S. technology companies, including leading AI developers, possess significant but fragmented intelligence on Chinese AI-enabled threats, but competitive concerns have prevented meaningful sharing. The AI-ISAC should enable developers to share indicators of compromise and emerging attack vectors without exposing proprietary information, providing the consolidated threat picture that feeds CISA’s public advisories and AI incident response frameworks. A priority workstream should be developing signatures of AI-driven exploitation and combining them with capability assessments to build an actionable picture of what systems and capabilities cyber actors find operationally useful.

The AI-ISAC should be modeled on existing sector-specific ISACs, with the Department of Justice issuing guidance clarifying that good-faith participation is protected under antitrust safe harbors.¹⁹⁷ Initial funding of \$10–15 million could be drawn from CISA’s Cybersecurity Division budget and additional congressional appropriations, supplemented by membership fees.

3

The Department of Commerce, in cooperation with the Departments of Defense, Energy, and State and U.S. industry, should establish security testing guidelines for U.S.-based entities that host, distribute, or provide managed access to AI models developed by foreign adversary entities.

NIST should establish guidelines for security testing and assurance of AI models developed by foreign adversary entities that are hosted, distributed, or provided managed access by U.S.-based entities, including testing for susceptibility to jailbreaking and prompt injection, ideological bias, presence of backdoors, and inadequate guardrails against misuse, among other areas. The guidelines should include public disclosure of testing results and assurance practices before hosting the model, enabling enterprise customers and government procurement officers to compare providers and thereby encourage more robust protections and analysis.

A secondary effort would be to develop best practices for U.S. entities to mitigate risks from AI models developed by foreign adversary entities, informed by aforementioned national security risk assessments from CAISI. Far from undermining their competitiveness, these guidelines would coordinate collective action and protect U.S. entities, including cloud service providers, from the greater reputational and legal exposure of becoming conduits for espionage or sabotage, which would harm enterprise customer trust far more than a narrowed model catalog.

4

The Department of Energy, leveraging infrastructure developed under the Genesis Mission, should establish a classified adversarial testing program for Chinese advanced AI systems at the national laboratories, in coordination with the Departments of Defense and Commerce and the Office of the Director of National Intelligence (ODNI).

The analytical gaps documented in this report are most acute in classified domains. While classification limits who can act on results, some of the most consequential questions about the capabilities of Chinese AI systems cannot be responsibly or feasibly tested in unclassified environments. How Chinese AI

systems perform when directed toward military planning, autonomous cyber exploitation, or biological weapons design, cannot be assessed through unclassified benchmarks alone.

The Genesis Mission's investment in an integrated AI platform at the labs provides the computational foundation without requiring new infrastructure or dependence on commercial cloud providers that could log sensitive queries.¹⁹⁸ The DoD and the ODNI should provide threat scenarios and operational requirements; Commerce should share CAISI evaluation data as baseline inputs. The classified program should be designed to produce declassifiable summaries and risk ratings that would in turn feed the public CAISI assessments recommended in this report, preserving operational detail while informing the broader policy response.

5

The Department of State should convene regular meetings with a core group of allies to share information on China's AI ecosystem and coordinate collective policy action.

Sustained, structured coordination among democratic allies is necessary to prevent Chinese AI systems from becoming the default digital infrastructure of the global economy. Existing multilateral efforts are inadequate to the task. The International Network for Advanced AI Measurement, Evaluation, and Science, which includes the United States and 10 other member governments, provides a useful forum for technical agencies to share evaluation methodologies.¹⁹⁹ But its membership includes countries that hedge between the United States and China, it does not operate at the policy level, and it lacks the mandate to coordinate restrictive actions or share intelligence. What is needed is a smaller, higher-level group.

The Department of State, led by the Bureau of Cyberspace and Digital Policy, should convene this group with support from the Departments of Commerce,

Defense, and Energy; the White House; and the ODNI. The group should begin with the United Kingdom, Japan, South Korea, Australia, Canada, Israel, the Netherlands, Germany, and Taiwan—allies with technical capacity and strategic alignment. Such a group is small enough to act and trusted enough to share sensitive information. Meetings should be held monthly at the assistant secretary-level or above, ensuring participants can make or directly inform decisions.

The group should establish three workstreams within four months. First, a common evaluation protocol for Chinese AI systems. Second, a coordinated disclosure framework that sequences the release of evaluation results and threat advisories. Third, a standing process for taking joint action against Chinese AI developers and systems of concern. Over time, this group could develop toward a formal institution for harmonizing technology protection measures.

One obvious potential pitfall is that this group becomes another coordination mechanism that meets regularly and accomplishes little. The workstreams recommended here are designed to force action, not just readouts and statements. If the group cannot produce measurable results within its first year, it should be dissolved.

6

The Department of Commerce, in cooperation with the Departments of Defense, Energy, and State and the ODNI, should publish semiannual reports on the state of China's AI ecosystem and supply chains to inform export controls, investment screening, and other restrictions.

Since 2017, nearly all actions to degrade the Chinese AI ecosystem have come from the executive branch, with little congressional involvement. Without regular reporting on China's AI ecosystem, Congress lacks the baseline to evaluate whether policy is working and

whether new authorities are needed. Regular assessments shared with Congress would support lawmakers in considering additional authorities, resourcing, and requirements for long-term strategic competition. The China AI Power Act, introduced in the House by Representative James Moynan on November 21, 2025, would fulfill this recommendation by requiring Commerce to submit annual reports on China's advanced AI capabilities and relevant supply chains.²⁰⁰ Semiannual rather than annual publication would better reflect the rapid pace of technological change.

1. Gregory C. Allen, *Understanding China's AI Strategy* (Center for a New American Security [CNAS], February 6, 2019), <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy>; Elsa B. Kania and Lorand Laskai, *Myths and Realities of China's Military-Civil Fusion Strategy* (CNAS, January 28, 2021), <https://www.cnas.org/publications/reports/myths-and-realities-of-chinas-military-civil-fusion-strategy>.
2. Allen, *Understanding China's AI Strategy*; Kania and Laskai, *Myths and Realities of China's Military-Civil Fusion Strategy*.
3. Sunny Cheung and Kai-shing Lau, "DeepSeek Use in PRC Military and Public Security Systems," *China Brief* 25, no. 20 (October 27, 2025), <https://jamestown.org/deepseek-use-in-prc-military-and-public-security-systems>; Kyle Miller et al., *The Use of Open Models in Research* (Center for Security and Emerging Technology [CSET], October 2025), <https://cset.georgetown.edu/publication/the-use-of-open-models-in-research>; Steven Feldstein, *The Global Expansion of AI Surveillance* (Carnegie Endowment for International Peace, September 17, 2019), <https://carnegieendowment.org/research/2019/09/the-global-expansion-of-ai-surveillance>.
4. *Evaluation of DeepSeek AI Models* (Center for AI Standards and Innovation [CAISI], National Institute of Standards and Technology, September 30, 2025), https://www.nist.gov/system/files/documents/2025/09/30/CAISI_Evaluation_of_DeepSeek_AI_Models.pdf.
5. Reuters, "Chinese AI Startup Zhipu Releases New Flagship Model GLM-5," February 11, 2026, <https://www.reuters.com/technology/chinas-ai-startup-zhipu-releases-new-flagship-model-glm-5-2026-02-11>; Jonathan Kemper, "Chinese AI Lab Zhipu Releases GLM-5 Under MIT License, Claims Parity with Top Western Models," *The Decoder*, February 12, 2026, <https://the-decoder.com/chinese-ai-lab-zhipu-releases-glm-5-under-mit-license-claims-parity-with-top-western-models>; Maxime Labonne, "GLM-5: China's First Public AI Company Ships a Frontier Model," *Hugging Face*, February 17, 2026, <https://huggingface.co/blog/mlabonne/glm-5>.
6. Kimi Team, "Kimi K2.5: Visual Agentic Intelligence," arXiv:2602.02276, February 2, 2026, <https://arxiv.org/abs/2602.02276>; Eduardo Baptista, "China's ByteDance Releases Doubao 2.0 AI Model for 'Agent Era,'" Reuters, February 14, 2026, <https://www.reuters.com/world/asia-pacific/chinas-bytedance-releases-doubao-20-ai-chatbot-2026-02-14/>; Eduardo Baptista, "Alibaba Unveils New Qwen3.5 Model for 'Agentic AI Era,'" Reuters, February 16, 2026, <https://www.reuters.com/world/china/alibaba-unveils-new-qwen35-model-agentic-ai-era-2026-02-16>.
7. Yuxuan Jia and Zichen Wang, "China's Spring Festival AI War," *Pekingology*, February 26, 2026, <https://www.pekingology.com/p/chinas-spring-festival-ai-war>; Irene Zhang and Nick Corvino, "Chinese AI Rings in the Year of the Horse: A Lunar New Year AI Roundup," *ChinaTalk*, February 18, 2026, <https://www.chinatalk.media/p/chinese-ai-rings-in-the-year-of-the>.
8. "OpenClaw Frenzy Drives China's Agentic AI Adoption, Raises Security Concerns," *Bloomberg*, March 12, 2026, <https://www.bloomberg.com/news/articles/2026-03-12/openclaw-frenzy-drives-china-s-agentic-ai-adoption-raises-security-concerns>; Luz Ding, "Alibaba Debuts OpenClaw App to Feed China's Agentic AI Addiction," *Bloomberg*, March 13, 2026, <https://www.bloomberg.com/news/articles/2026-03-13/alibaba-debuts-openclaw-app-to-feed-china-s-agentic-ai-addiction>; Josh Xiao and Nectar Gan, "China's OpenClaw Frenzy Test Xi's Approach to Regulate AI," *Bloomberg*, March 12, 2026, <https://www.bloomberg.com/news/articles/2026-03-12/china-s-openclaw-frenzy-tests-xi-s-approach-to-regulate-ai>.
9. *15th Five-Year Plan for Economic and Social Development of the People's Republic of China (2026–2030): Draft Outline* (NPC Observer, March 2026), https://npcobserver.com/wp-content/uploads/2026/03/15th-Five-Year-Plan-Draft_NON-FINAL.pdf.
10. *Winning the Race: America's AI Action Plan* (White House Office of Science and Technology Policy, July 23, 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.
11. Maggie Hassan and Joni Ernst, "Senators Hassan, Ernst Sound Alarm on Chinese AI-Enabled Hackers," press release, December 3, 2025, <https://www.hassan.senate.gov/news/press-releases/senators-hassan-ernst-sound-alarm-on-chinese-ai-enabled-hackers>; *Evaluation of DeepSeek AI Models*.
12. Michael Kratsios, "Remarks at the Security Council's Open Debate on Artificial Intelligence and International Peace and Security," United States Mission to the United Nations, September 24, 2025, <https://usun.usmission.gov/remarks-at-the-security-councils-open-debate-on-artificial-intelligence-and-international-peace-and-security>.
13. Liza Tobin, "China's Brute Force Economics: Waking Up from the Dream of a Level Playing Field," *Texas National Security Review* 6, no. 1 (2022/2023): 81–98, <https://tnsr.org/2022/12/chinas-brute-force-economics-waking-up-from-the-dream-of-a-level-playing-field>.
14. Jinghan Zeng et al., "Securing China's Core Interests: The State of the Debate in China," *International Affairs* 91, no. 2 (2015): 245–266, https://www.chathamhouse.org/sites/default/files/field/field_document/INTA91_2_03_Zeng_Xiao_Breslin.pdf.
15. Steve Tsang, "Party-State Realism: A Framework for Understanding China's Approach to Foreign Policy," *Journal of Contemporary China* 29, no. 122 (2020): 304, <https://www.tandfonline.com/doi/abs/10.1080/10670564.2019.1637562>; Susan L. Shirk, "China in Xi's New Era: The Return to Personalistic Rule," *Journal of Democracy* 29, no. 2 (2018): 22–36, <https://www.journalofdemocracy.org/articles/china-in-xis-new-era-the-return-to-personalistic-rule/>; and Margaret M. Pearson et al., "China's Party-State Capitalism and International Backlash: From Interdependence to Insecurity," *International Security* 47, no. 2 (2022): 135, https://doi.org/10.1162/isec_a_00447.
16. Rush Doshi, *The Long Game: China's Grand Strategy to Displace American Order* (Oxford University Press, 2021), 261–264; Elizabeth C. Economy, *The Third Revolution: Xi Jinping and the New Chinese State* (Oxford University Press, 2018); Bethany Allen, *Beijing Rules: How China Weaponized Its Economy to Confront the World* (Harper, 2023); and Michael J. Mazarr

- et al., *China and the International Order* (RAND Corporation, 2018), https://www.rand.org/pubs/research_reports/RR2423.html.
17. National Counterintelligence and Security Center, “U.S. Business Risk: People’s Republic of China (PRC) Laws Expand Beijing’s Oversight of Foreign and Domestic Companies,” *Safeguarding Our Future* (bulletin), June 20, 2023, https://www.dni.gov/files/NCSC/documents/SafeguardingOurFuture/FINAL_NCSC_SOF_Bulletin_PRC_Laws.pdf.
 18. “PRC National Intelligence Law (as Amended in 2018),” China Law Translate, accessed March 31, 2026, <https://www.chinalawtranslate.com/en/national-intelligence-law-of-the-prc-2017>.
 19. Rogier Creemers et al., trans., “Translation: Cybersecurity Law of the People’s Republic of China (Effective June 1, 2017),” *DigiChina*, Stanford University, June 29, 2018, <https://digichina.stanford.edu/work/translation-cybersecurity-law-of-the-peoples-republic-of-china-effective-june-1-2017>.
 20. Deng Xiaoping, “Speech at the Opening Ceremony of the National Conference On Science,” Marxists Internet Archive, March 18, 1978, <https://www.marxists.org/reference/archive/deng-xiaoping/1978/30.htm>; Xi Jinping, “Striving to Build a Country Strong in Science and Technology,” *Qiushi*, July 15, 2025, https://en.qstheory.cn/2025-07/15/c_1109429.htm; Manoj Kewalramani, “He Lifeng on China’s Approach to Developing New-Quality Productive Forces,” *Tracking People’s Daily*, November 11, 2025, <https://trackingpeoplesdaily.substack.com/p/he-lifeng-on-chinas-approach-to-developing>.
 21. Xi, “Striving to Build a Country Strong in Science and Technology”; Arthur R. Kroeber, “Unleashing ‘New Quality Productive Forces’: China’s Strategy for Technology-Led Growth,” Brookings Institution, June 4, 2024, <https://www.brookings.edu/articles/unleashing-new-quality-productive-forces-chinas-strategy-for-technology-led-growth>; Xinhua News Agency, “授权发布 | 习近平：在全国科技大会、国家科学技术奖励大会、两院院士大会上的讲话 [Authorized Publication | Xi Jinping: Speech at the National Science and Technology Conference, the National Science and Technology Awards Conference, and the Academician Conference of the Two Academies],” Xinhua, June 24, 2024, <https://www.news.cn/politics/leaders/20240624/16741a201e564d8d8775ffb1450ecf29/c.html>.
 22. Jeffrey Ding, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition* (Princeton University Press, 2024); Paul Kennedy, *The Rise and Fall of the Great Powers: Economic Change and Military Conflict from 1500 to 2000* (Random House, 1987).
 23. Josh Chin and Liza Lin, *Surveillance State: Inside China’s Quest to Launch a New Era of Social Control* (St. Martin’s Press, 2022); Minxin Pei, *The Sentinel State: Surveillance and the Survival of Dictatorship in China* (Harvard University Press, 2024), 213–237; Jennifer Lind, *Autocracy 2.0: How China’s Rise Reinvented Tyranny* (Cornell University Press, 2025), 60–62.
 24. Joel Wuthnow and Phillip C. Saunders, *China’s Quest for Military Supremacy* (Polity, 2025); Fiona S. Cunningham, *Under the Nuclear Shadow: China’s Information-Age Weapons in International Security* (Princeton University Press, 2025), 185–241; Jordan Schneider and Irene Zhang, “China + AI = Military Advantage? Plus: DC Meetup!,” *ChinaTalk*, January 20, 2023, <https://www.chinatalk.media/p/china-ai-military-advantage-plus>.
 25. Maarten Buyl et al., “Large Language Models Reflect the Ideology of Their Creators,” *npj Artificial Intelligence* 2, no. 1 (2026): 7, <https://www.nature.com/articles/s44387-025-00048-0>; Jorge Perez, “Tokenising Culture: Causes and Consequences of Cultural Misalignment in Large Language Models,” *Ada Lovelace Institute*, June 19, 2025, <https://www.adalovelaceinstitute.org/blog/cultural-misalignment-llms>; Emanuel Z. Fenech-Borg et al., “The Cultural Gene of Large Language Models: A Study on the Impact of Cross-Corpus Training on Model Values and Biases,” arXiv:2508.12411, December 29, 2025, <https://arxiv.org/abs/2508.12411>.
 26. Nathan Beauchamp-Mustafaga, *Exploring the Implications of Generative AI for Chinese Military Cyber-Enabled Influence Operations: Chinese Military Strategies, Capabilities, and Intent* (RAND Corporation, February 1, 2024), <https://www.rand.org/pubs/testimonies/CTA3191-1.html>; “Document 9: A ChinaFile Translation: How Much Is a Hardline Party Directive Shaping China’s Current Political Climate?” *ChinaFile*, November 8, 2013, <https://www.chinafile.com/document-9-chinafile-translation>.
 27. Jennifer Lind, “China’s Smart Authoritarianism: How the CCP Balances Control and Innovation,” *Foreign Affairs*, February 10, 2026, <https://www.foreignaffairs.com/china/chinas-smart-authoritarianism>.
 28. “Military-Civil Fusion,” U.S. Department of State, accessed March 31, 2026, <https://2017-2021.state.gov/military-civil-fusion>; Kania and Laskai, *Myths and Realities of China’s Military-Civil Fusion Strategy*; Jordan Schneider, “Choking Off China’s AI Access: Exploring the Impact of the Commerce Department’s New Export Controls,” *ChinaTalk*, October 12, 2022, <https://www.chinatalk.media/p/choking-off-chinas-ai-access>.
 29. Graham Webster et al., trans., “Full Translation: China’s ‘New Generation Artificial Intelligence Development Plan’ (2017),” *DigiChina*, Stanford University, August 1, 2017, <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017>.
 30. “国务院关于深入实施‘人工智能+’行动的意见 [Opinions of the State Council on the in-Depth Implementation of the ‘Artificial Intelligence+’ Action],” State Council of the People’s Republic of China, August 21, 2025, https://www.gov.cn/zhengce/content/202508/content_7037861.htm; Matt Sheehan, “China’s Big AI Diffusion Plan Is Here. Will It Work? A Bearish and Bullish Case for the Major ‘AI+’ Initiative,” *Interconnected*, September 9, 2025, <https://mattsheehan.substack.com/p/chinas-big-ai-diffusion-plan-is-here>; “China Releases ‘AI Plus’ Policy: A Brief Analysis,” *Geopolitechs*, August 26, 2025, <https://www.geopolitechs.org/p/china-releases-ai-plus-policy-a-brief>; Ruby Osman, “Preparing for ‘Changes Unseen in a Century’: What to Expect From China’s New Five-Year Plan,” *Tony Blair Institute for Global Change*, February 18, 2026, <https://institute.global/insights/geopolitics-and-security/what-to-expect-from-chinas-new-five-year-plan>; Celeste Li, “AI Management in China’s 15th Five-Year Plan,” *Lambheart*, March 15, 2026, <https://lambheart>.

substack.com/p/ai-management-in-chinas-15th-five; “AI in the Five-Year Plan; Regulating AI Agents; Qwen and the Two Sessions,” *Lingua Sinica*, March 6, 2026, <https://linguasinica.substack.com/p/ai-in-the-five-year-plan-regulating>.

31. Vivek Chilukuri and Ruby Scanlon, *Countering the Digital Silk Road* (CNAS, October 15, 2025), <https://www.cnas.org/publications/reports/countering-the-digital-silk-road>; Martin Beraja et al., “Exporting the Surveillance State via Trade in AI,” Brookings Institution, January 12, 2023, <https://www.brookings.edu/articles/exporting-the-surveillance-state-via-trade-in-ai>.
32. Rogier Creemers et al., trans., “Translation: Internet Information Service Algorithmic Recommendation Management Provisions (Effective March 1, 2022),” *DigiChina*, Stanford University, January 10, 2022, <https://digichina.stanford.edu/work/translation-internet-information-service-algorithmic-recommendation-management-provisions-effective-march-1-2022>; “Translation: Interim Measures for the Management of Generative Artificial Intelligence Services,” *China Law Translate*, July 13, 2023, <https://www.chinalawtranslate.com/en/generative-ai-interim>.
33. Matt Sheehan and Scott Singer, “How China Views AI Risks and What to do About Them,” Carnegie Endowment for International Peace, October 16, 2025, <https://carnegieendowment.org/research/2025/10/how-china-views-ai-risks-and-what-to-do-about-them>; “人工智能安全治理框架 2.0 [AI Safety Governance Framework 2.0],” *Cyberspace Administration of China*, September 15, 2025, https://www.cac.gov.cn/2025-09/15/c_1759653448369123.htm; “China Releases Upgraded AI Safety Governance Framework to Tackle Emerging AI Risks,” *Geopolitechs*, September 15, 2025, <https://www.geopolitechs.org/p/china-releases-upgraded-ai-safety>.
34. Yi-Ling Liu, “Thousands of Companies Are Driving China’s AI Boom. A Government Registry Tracks Them All,” *Wired*, January 20, 2026, <https://www.wired.com/story/china-ai-boom-algorithm-registry>; Nicholas Welch, “SB 1047 with Socialist Characteristics: China’s Algorithm Registry in the LLM Era,” *ChinaTalk*, November 14, 2024, <https://www.chinatalk.media/p/sb-1047-with-socialist-characteristics>; Oxford China Policy Lab (OCPL) and Zilan Qian, “Expert Insight: China’s AI Services Registry System, A Complete Guide,” OCPL, January 23, 2026, <https://ocpl.substack.com/p/expert-insight-chinas-ai-services>; Nick Corvino, “Making Money in Chinese AI Safety,” *ChinaTalk*, March 12, 2026, <https://www.chinatalk.media/p/the-business-behind-chinese-ai-safety>.
35. Matt Sheehan, “China’s AI Companion Reg & the Growing Role of Technical Standards,” *Interconnected*, March 3, 2026, <https://mattsheehan.substack.com/p/chinas-ai-companion-reg-and-the-growing>; Nick Corvino, “Making Money in Chinese AI Safety: Compliance-as-a-Service,” *ChinaTalk*, March 12, 2026, <https://www.chinatalk.media/p/the-business-behind-chinese-ai-safety>.
36. Elizabeth C. Economy, *The World According to China* (Polity Press, 2021), 171–175; Daniel F. Runde and Austin Hardman, “Great Power Competition in the Multilateral System,” *Center for Strategic and International Studies* (CSIS), October 23, 2024, <https://www.csis.org/analysis/great-power-competition-multilateral-system>; Sebastian Haug et al., “Power Shifts in International Organisations: China at the United Nations,” *Global Policy* 15, Suppl. 2 (2024): 5–17, <https://onlinelibrary.wiley.com/doi/10.1111/1758-5899.13368>; Joel Wuthnow et al., “Diverse Multilateralism: Four Strategies in China’s Multilateral Diplomacy,” *Journal of Chinese Political Science* 17, no. 3 (2012): 269–290, https://www.researchgate.net/publication/257680021_Diverse_Multilateralism_Four_Strategies_in_China%27s_Multilateral_Diplomacy.
37. “Global AI Governance Initiative,” Ministry of Foreign Affairs of the People’s Republic of China, October 20, 2023, https://www.mfa.gov.cn/mfa_eng/zy/gb/202405/t20240531_11367503.html; “Global AI Governance Action Plan,” Ministry of Foreign Affairs of the People’s Republic of China, July 26, 2025, https://www.fmprc.gov.cn/mfa_eng/xw/zyxw/202507/t20250729_11679232.html.
38. “Global AI Governance Action Plan.”
39. Ministry of Industry and Information Technology et al., *国家人工智能产业综合标准化体系建设指南(2024版) [Guidelines for the Construction of a Comprehensive Standardization System for the National Artificial Intelligence Industry (2024 Edition)]* (CSET, October 1, 2025), <https://cset.georgetown.edu/publication/china-ai-standardization-guidelines-2024>.
40. “Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification,” 87 Fed. Reg. 62186 (October 13, 2022), <https://www.federalregister.gov/documents/2022/10/13/2022-21658/implementation-of-additional-export-controls-certain-advanced-computing-and-semiconductor>; “Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections,” 88 Fed. Reg. 73458 (October 25, 2023), <https://www.federalregister.gov/documents/2023/10/25/2023-23055/implementation-of-additional-export-controls-certain-advanced-computing-items-supercomputer-and>.
41. U.S. Department of Commerce, “The Department of Commerce Announces American AI Exports Program Implementation,” press release, October 21, 2025, <https://www.trade.gov/press-release/department-commerce-announces-american-ai-exports-program-implementation>; “Pax Silica,” U.S. Department of State, accessed March 31, 2026, <https://www.state.gov/pax-silica>.
42. Helen Toner, “What Are Generative AI, Large Language Models, and Foundation Models?” CSET, May 12, 2023, <https://cset.georgetown.edu/article/what-are-generative-ai-large-language-models-and-foundation-models>; Thomas Woodside and Helen Toner, “Multimodality, Tool Use, and Autonomous Agents: Large Language Models Explained, Part 3,” CSET, March 8, 2024, <https://cset.georgetown.edu/article/multimodality-tool-use-and-autonomous-agents>; Ryan Lopopolo, “Harness Engineering: Leveraging Codex in an Agent-First World,” *OpenAI*, February 11, 2026, <https://openai.com/index/harness-engineering>.
43. *Hiroshima Process International Code of Conduct for Organi-*

- zations Developing Advanced AI Systems (Ministry of Foreign Affairs of Japan, October 30, 2023), <https://www.mofa.go.jp/files/100573473.pdf>; “General-Purpose AI Models in the AI Act – Questions & Answers,” European Commission, accessed March 31, 2026, <https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers>; Vibhu Mishra, “General Assembly Adopts Landmark Resolution on Artificial Intelligence,” UN News, March 21, 2024, <https://news.un.org/en/story/2024/03/1147831>; Yoshua Bengio et al., *International AI Safety Report 2025* (International AI Safety Report, January 29, 2025), <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025>.
44. Sunny Cheung and Kai-shing Lau, “DeepSeek Use in PRC Military and Public Security Systems,” *China Brief* 25, no. 20 (2025), <https://jamestown.org/deepseek-use-in-prc-military-and-public-security-systems>; Haye Kesteloo, “China’s Military Deploys Cost-Efficient DeepSeek AI Across Drone Swarms and Robot Dogs,” *DroneXL*, October 28, 2025, <https://dronexl.co/2025/10/28/china-military-deepseek-ai-drone-swarms-robot-dogs>; Yong Jian, “DeepSeek Is Now the Brain of Chinese State-Owned Firms,” *Asia Times*, February 28, 2025, <https://asiatimes.com/2025/02/deepseek-is-now-the-brain-of-chinese-state-owned-firms>; Kinling Lo, “China’s AI Frenzy: DeepSeek Is Already Everywhere – Cars, Phones, Even Hospitals,” *Rest of World*, March 13, 2025, <https://restofworld.org/2025/china-embeds-deepseek-ai-in-everything>; Samuel Wade, “Translations: DeepSeek’s ‘Outstanding Results in the Field’ of Public Security and Public Opinion Response,” *China Digital Times*, April 9, 2025, <https://chinadigitaltimes.net/2025/04/translations-deepseeks-outstanding-results-in-the-field-of-public-security>.
 45. Kaytie Ward, “What Is an Open-Weights Model?,” *AI21*, May 27, 2025, <https://www.ai21.com/glossary/foundational-llm/open-weights-model>; Kyle Miller et al., “The Use of Open Models in Research,” *CSET*, October 2025, <https://cset.georgetown.edu/publication/the-use-of-open-models-in-research>.
 46. Sam Eifling, “China’s Biggest AI Model Is Challenging American Dominance,” *Rest of World*, September 23, 2024, <https://restofworld.org/2024/alibaba-qwen-ai-model>; “MediaTek Dimensity 9400 Fully Supports Qwen3’s Advanced AI Features,” *MediaTek* (blog), May 14, 2025, <https://www.mediatek.com/tek-talk-blogs/mediatek-dimensity-9400-fully-supports-qwen3s-advanced-ai-features>; Christopher Ort, “Qwen 3.5 Small: Alibaba’s On-Device AI Innovation,” *i10x*, March 3, 2026, <https://i10x.ai/news/qwen-3-5-small-alibaba-on-device-ai-models>.
 47. Chris McGuire, “China’s AI Chip Deficit: Why Huawei Can’t Catch Nvidia and U.S. Export Controls Should Remain,” *Council on Foreign Relations*, December 15, 2025, <https://www.cfr.org/articles/chinas-ai-chip-deficit-why-huawei-cant-catch-nvidia-and-us-export-controls-should-remain>; Saif M. Khan, “Securing Semiconductor Supply Chains,” *CSET*, January 2021, <https://cset.georgetown.edu/publication/securing-semiconductor-supply-chains/>; Kari Briski, “How Scaling Laws Drive Smarter, More Powerful AI,” *NVIDIA* (blog), February 12, 2025, <https://blogs.nvidia.com/blog/ai-scaling-laws>.
 48. Raffaele Huang and Tracy Qu, “Chinese AI Developers Say They Can’t Beat America Without Better Chips,” *The Wall Street Journal*, January 15, 2026, <https://www.wsj.com/tech/ai/china-ai-race-us-chips-9e74b957>; “[AINews] Z.ai GLM-5: New SOTA Open Weights LLM,” *Latent Space*, February 12, 2026, <https://www.latent.space/p/ainews-zai-glm-5-new-sota-open-weights>; Z.ai (@Zai_org), “GLM-5 is coming to Coding Plan Pro users within one week, and we’re working to bring it to everyone after that,” *X*, February 11, 2026, https://x.com/Zai_org/status/2021656633320018365.
 49. Saif Khan et al., “Should the US Sell Hopper Chips to China?,” *Institute for Progress*, December 7, 2025, <https://ifp.org/should-the-us-sell-hopper-chips-to-china/>.
 50. Lisa Klaassen and Broderick McDonald, “America Is Running the Wrong AI Race,” *The National Interest*, January 27, 2026, <https://nationalinterest.org/blog/techland/america-is-running-the-wrong-ai-race>.
 51. Tracy Qu, “China’s Alibaba Links Qwen AI App to Vast Consumer Ecosystem,” *The Wall Street Journal*, January 14, 2026, <https://www.wsj.com/tech/chinas-alibaba-links-qwen-ai-app-to-vast-consumer-ecosystem-17b4f942>.
 52. Jordan Schneider et al., “The Z.ai Playbook: Zixuan Li on the Chinese AI ecosystem,” *ChinaTalk*, November 21, 2025, <https://www.chinatalk.media/p/the-zai-playbook>; Irene Zhang et al., “The All-Star Chinese AI Conversation of 2026,” *ChinaTalk*, January 13, 2026, <https://www.chinatalk.media/p/the-all-star-chinese-ai-conversation>.
 53. Zilan Qian, “China’s AI Landscape: A Free-for-All, Not a Central Plan,” *ChinaTalk*, January 30, 2026, <https://www.chinatalk.media/p/chinas-ai-landscape-a-free-for-all>.
 54. Florian Brand and Nathan Lambert, “2025 Open Models Year in Review,” *Interconnects*, December 14, 2025, <https://www.interconnects.ai/p/2025-open-models-year-in-review>.
 55. “LLM Leaderboard – Comparison of over 100 AI models from OpenAI, Google, DeepSeek & others,” *Artificial Analysis*, accessed March 31, 2026, <https://artificialanalysis.ai/leaderboards/models>.
 56. Jeffrey Ding, “ChinAI #345: A Three-Way Race for China’s AI Super-App,” *ChinAI Newsletter*, February 2, 2026, <https://chinai.substack.com/p/chinai-345-a-three-way-race-for-chinas>.
 57. Kevin Xu, “No Business Model: DeepSeek’s Enduring Advantage,” *Interconnected*, January 16, 2026, <https://interconnected.blog/no-business-model-deepseeks-enduring-advantage>.
 58. Liu Peilin, “China’s Zhipu AI Jumps in Hong Kong Debut,” *Caixin Global*, January 8, 2026, <https://www.caixinglobal.com/2026-01-08/chinas-zhipu-ai-jumps-in-hong-kong-debut-102401610.html>; Schneider et al., “The Z.ai Playbook.”
 59. *Addition of Entities to and Revision of Entry on the Entity List* (Bureau of Industry and Security, Department of Commerce, January 16, 2025), <https://public-inspection.federalregister.gov/2025-00704.pdf>; Knowledge Atlas Technology Joint

- Stock Company Limited (Zhipu AI), *Global Offering Prospectus* (HKEX, December 30, 2025), <https://www1.hkexnews.hk/listedco/listconews/sehk/2025/1230/2025123000017.pdf>; Jiang Jiang, “Zhipu CEO Zhang Peng on Taking China’s First LLM Company Public,” *Ginger River Review*, January 11, 2026, <https://www.gingerriver.com/p/zhipu-ceo-zhang-peng-on-taking-chinas>.
60. Liu Peilin et al., “In Depth: Global Capital, Chinese AI Converge in Hong Kong,” *Caixin Global*, February 6, 2026, <https://www.caixinglobal.com/2026-02-06/in-depth-global-capital-chinese-ai-converge-in-hong-kong-102412145.html>.
 61. Nathan Lambert, “Open Models in Perpetual Catch-Up,” *Interconnects*, February 17, 2026, <https://www.interconnects.ai/p/open-models-in-perpetual-catch-up>; “The End of SWE-Bench Verified – Mia Glaese and Olivia Watkins, OpenAI Frontier Evals & Human Data,” *Latent Space*, February 23, 2026, <https://www.latent.space/p/swe-bench-dead>.
 62. Jacob Stokes, Alexander Sullivan, and Noah Greene, *U.S.-China Competition and Military AI: How Washington Can Manage Strategic Risks amid Rivalry with Beijing* (CNAS, July 25, 2023), <https://www.cnas.org/publications/reports/u-s-china-competition-and-military-ai>.
 63. Dan Tadross and Jared Jonker, *Agentic Warfare* (Scale AI, January 2026), <https://scale.com/agentic-warfare>; Zachary Burdette et al., *How Artificial Intelligence Could Reshape Four Essential Competitions in Future Warfare* (RAND Corporation, January 22, 2026), https://www.rand.org/pubs/research_reports/RRA4316-1.html.
 64. William N. Caballero and Phillip R. Jenkins, “On Large Language Models in National Security Applications,” arXiv:2407.03453, July 3, 2024, <https://arxiv.org/abs/2407.03453>; Richard Farnell and Kira Coffey, “AI’s New Frontier in War Planning: How AI Agents Can Revolutionize Military Decision-Making,” *Belfer Center for Science and International Affairs*, October 11, 2024, <https://www.belfercenter.org/research-analysis/ais-new-frontier-war-planning-how-ai-agents-can-revolutionize-military-decision>; Vincenzo Gallitelli, “Artificial Intelligence-Enabled Military Decision-Making Process,” *Journal of Advanced Military Studies* 16, no. 2 (2025), <https://www.usmcu.edu/Outreach/Marine-Corps-University-Press/MCU-Journal/JAMS-vol-16-no-2/Artificial-Intelligence-Enabled-Military-Decision-Making-Process>.
 65. Ritwik Gupta, “LLMs, Autonomous Weapons, and AI Research,” *ritwikgupta.me* (blog), March 2, 2026, <https://ritwikgupta.me/blog/posts/2026-03-02-llms-autonomous-weapons-ai-research/>.
 66. Tara Copp et al., “Anthropic’s AI Tool Claude Central to U.S. Campaign in Iran, amid a Bitter Feud,” *The Washington Post*, March 4, 2026, <https://www.washingtonpost.com/technology/2026/03/04/anthropic-ai-iran-campaign/>; Katrina Manson, “Iran War Tests Project Maven, U.S. AI War Strategy,” *Bloomberg*, March 12, 2026, <https://www.bloomberg.com/news/features/2026-03-12/iran-war-tests-project-maven-us-ai-war-strategy>.
 67. Emelia Probasco et al., *China’s Military AI Wish List* (CSET, February 2026), <https://cset.georgetown.edu/publication/chinas-military-ai-wish-list>.
 68. Cole McFaul et al., *Pulling Back the Curtain on China’s Military-Civil Fusion* (CSET, September 2025), <https://cset.georgetown.edu/publication/pulling-back-the-curtain-on-chinas-military-civil-fusion>; Insikt Group, “Artificial Eyes: Generative AI in China’s Military Intelligence,” *Recorded Future*, June 17, 2025, <https://www.recordedfuture.com/research/artificial-eyes-generative-ai-chinas-military-intelligence>; Probasco et al., *China’s Military AI Wish List*.
 69. McFaul et al., *Pulling Back the Curtain on China’s Military-Civil Fusion*; Jessica C. Liao and Joshua Arostegui, “Can the 15th Five-Year Plan Fix the People’s Liberation Army’s Procurement Bottlenecks?” *War on the Rocks*, January 14, 2026, <https://warontherocks.com/2026/01/can-the-15th-five-year-plan-fix-the-peoples-liberation-armys-procurement-bottlenecks/>.
 70. Sunny Cheung and Kai-shing Lau, “DeepSeek Use in PRC Military and Public Security Systems,” *China Brief* 25, no. 20 (2025), <https://jamestown.org/deepseek-use-in-prc-military-and-public-security-systems/>; Eduardo Baptista and Fanny Potkin, “How China Could Use DeepSeek in an Era of War,” *Reuters*, October 27, 2025, <https://www.reuters.com/world/asia-pacific/robot-dogs-ai-drone-swarms-how-china-could-use-deepseek-an-era-war-2025-10-27/>.
 71. Alex Colville, “The Central Military Commission on AI Use in the Army,” *Lingua Sinica*, January 29, 2026, <https://linguasinica.substack.com/p/the-central-military-commission-on>; Wang Gongjin and Luo Yu, “人工智能在军事领域的主要应用与前景展望 [The Main Applications and Prospects of Artificial Intelligence in the Military Field],” *国防科技 [National Defense Technology]* 46, no. 4 (August 2025), <https://drive.google.com/file/d/1Vym-rPMXufPTkmO-Vi1yfdxhY3VlLhik/view>.
 72. Kimi Team, “Kimi K2: Open Agentic Intelligence,” arXiv:2507.20534, February 23, 2026, <https://arxiv.org/abs/2507.20534>; “Kimi K2.5: Visual Agentic Intelligence,” *Kimi*, January 27, 2026, <https://www.kimi.com/blog/kimi-k2-5>; DeepSeek-AI, “DeepSeek-V3.2: Pushing the Frontier of Open Large Language Models,” arXiv:2512.02556, December 2, 2025, <https://arxiv.org/abs/2512.02556>; “Why SWE-Bench Verified No Longer Measures Frontier Coding Capabilities,” *OpenAI*, February 23, 2026, <https://openai.com/index/why-we-no-longer-evaluate-swe-bench-verified/>.
 73. Carl Franzen, “DeepSeek-V3.1-Terminus Launches with Improved Agentic Tool Use and Reduced Language Mixing Errors,” *VentureBeat*, September 22, 2025, <https://venturebeat.com/ai/deepseek-v3-1-terminus-launches-with-improved-agentic-tool-use-and-reduced/>; “Kimi K2.5: Visual Agentic Intelligence”; “GLM-5,” *Z.AI Developer Documentation*, accessed April 3, 2026, <https://docs.z.ai/guides/llm/glm-5>.
 74. Copp et al., “Anthropic’s AI Tool Claude Central to U.S. Campaign in Iran, amid a Bitter Feud”; Gupta, “LLMs, Autonomous Weapons, and AI Research”; Manson, “Iran War Tests Project Maven, U.S. AI War Strategy.”

75. Gupta, "LLMs, Autonomous Weapons, and AI Research."
76. Gupta, "LLMs, Autonomous Weapons, and AI Research"; Michael C. Horowitz, "Autonomous Weapon Systems: No Human-in-the-Loop Required, and Other Myths Dispelled," War on the Rocks, May 22, 2025, <https://warontherocks.com/2025/05/autonomous-weapon-systems-no-human-in-the-loop-required-and-other-myths-dispelled/>.
77. Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (W.W. Norton, 2018).
78. Horowitz, "Autonomous Weapon Systems."
79. Jiedong Lang et al., "A Comprehensive Study on Quantization Techniques for Large Language Models," arXiv:2411.02530, October 30, 2024, <https://arxiv.org/abs/2411.02530>.
80. Abdolmaged Alkhulaifi, Fahad Alsahli, and Irfan Ahmad, "Knowledge Distillation in Deep Learning and its Applications," arXiv:2007.09029, July 17, 2020, <https://arxiv.org/abs/2007.09029>.
81. Caroline Meinhardt et al., trans., "Beyond DeepSeek: China's Diverse Open-Weight AI Ecosystem and Its Policy Implications," Stanford Human-Centered Artificial Intelligence/DigiChina, December 16, 2025, <https://hai.stanford.edu/assets/files/hai-digichina-issue-brief-beyond-deepseek-chinas-diverse-open-weight-ai-ecosystem-policy-implications.pdf>; Yunna Li et al., "Analysis of Key Technologies in Edge Computing and Practical Applications in U.S. Military Operations," in *Proceedings of the 2025 8th International Conference on Computer Information Science and Artificial Intelligence* (Association for Computing Machinery, December 19, 2025), 1183–1190, <https://dl.acm.org/doi/10.1145/3773365.3773552>.
82. Caleb Withers, *Tipping the Scales: Emerging AI Capabilities and the Cyber Offense-Defense Balance* (CNAS, September 23, 2025), <https://www.cnas.org/publications/reports/tipping-the-scales>; Sean Heelan, "On the Coming Industrialisation of Exploit Generation with LLMs," *Sean Heelan's Blog*, January 18, 2026, <https://sean.heelan.io/2026/01/18/on-the-coming-industrialisation-of-exploit-generation-with-llms/>.
83. Anthropic Threat Intelligence, "Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign," Anthropic, November 17, 2025, <https://assets.anthropic.com/m/ec212e6566a0d47/original/Disrupting-the-first-reported-AI-orchestrated-cyber-espionage-campaign.pdf>.
84. Andy K. Zhang et al., "Cybench: A Framework for Evaluating Cybersecurity Capabilities and Risks of Language Models," arXiv:2408.08926, April 12, 2025, <https://arxiv.org/abs/2408.08926>; "Evaluations: Cyber Offense," Frontier AI Risk Monitoring Platform, accessed April 3, 2026, <https://airiskmonitor.net/domain/cyber/evaluation>; Nathan Lambert, "Opus 4.6, Codex 5.3, and the Post-Benchmark Era," Interconnects, February 9, 2026, <https://www.interconnects.ai/p/opus-46-vs-codex-53>.
85. Frontier AI Risk Monitoring Platform, "Evaluations: Cyber Offense."
86. Kathleen C. Fraser et al., "Fine-Tuning Lowers Safety and Disrupts Evaluation Consistency," arXiv:2506.17209, June 20, 2025, <https://arxiv.org/abs/2506.17209>.
87. Frontier AI Risk Monitoring Platform, "Risk Analysis: Cyber Offense," airiskmonitor.net, accessed April 3, 2026, <https://airiskmonitor.net/domain/cyber/risk>.
88. Zhang et al., "Cybench"; Frontier AI Risk Monitoring Platform, "Evaluations: Cyber Offense"; *Evaluation of DeepSeek AI Models*; "CAISI Evaluation of Kimi K2 Thinking," National Institute of Standards and Technology (NIST), January 9, 2026, <https://www.nist.gov/news-events/news/2025/12/caisi-evaluation-kimi-k2-thinking>; *System Card: Claude Opus 4 & Claude Sonnet 4* (Anthropic, May 2025), <https://www-cdn.anthropic.com/6be99a52cb68eb70eb9572b4cafad13df32ed995.pdf>; *System Card Addendum: Claude Opus 4.1* (Anthropic, August 2025), <https://www-cdn.anthropic.com/9fa30625273bafdf5af82c93719d7ca606485a16.pdf>; *System Card: Claude Sonnet 4.5* (Anthropic, September 2025), <https://assets.anthropic.com/m/12f214efcc2f457a/original/Claude-Sonnet-4-5-System-Card.pdf>; *System Card: Claude Opus 4.5* (Anthropic, November 2025), <https://assets.anthropic.com/m/64823ba7485345a7/Claude-Opus-4-5-System-Card.pdf>; *System Card: Claude Opus 4.6* (Anthropic, February 2026), <https://www-cdn.anthropic.com/Odd865075ad3132672ee0ab40b05a53f14cf5288.pdf>; *Grok 4 Model Card* (xAI, August 20, 2025), <https://data.x.ai/2025-08-20-grok-4-model-card.pdf>; *Grok 4 Fast Model Card* (xAI, September 19, 2025), <https://data.x.ai/2025-09-19-grok-4-fast-model-card.pdf>; *Grok 4.1 Model Card* (xAI, November 17, 2025), <https://data.x.ai/2025-11-17-grok-4-1-model-card.pdf>.
89. Zhang et al., "Cybench"; Frontier AI Risk Monitoring Platform, "Evaluations: Cyber Offense"; *Evaluation of DeepSeek AI Models*; "CAISI Evaluation of Kimi K2 Thinking"; Anthropic, *Claude Opus 4, Sonnet 4*; Anthropic, *Claude Opus 4.1 Addendum*; *Claude Sonnet 4.5 System Card* (Anthropic, September 2025), <https://assets.anthropic.com/m/12f214efcc2f457a/original/Claude-Sonnet-4-5-System-Card.pdf>; Anthropic, *Claude Opus 4.5*; Anthropic, *Claude Opus 4.6*"; xAI, *Grok 4 Model Card*; xAI, *Grok 4 Fast Model*; xAI, *Grok 4.1 Model Card*; Zheng-Xin Yong et al., "An Independent Safety Evaluation of Kimi K2.5," arXiv:2604.03121v1, April 3, 2026, <https://arxiv.org/pdf/2604.03121>.
90. Roger Brent and Greg McKelvey, Jr., "Contemporary Foundation AI Models Increase Biological Weapons Risk," RAND Corporation, December 31, 2025, <https://www.rand.org/pubs/perspectives/PEA3853-1.html>.
91. Kyle Brady et al., "Bridging the Digital to Physical Divide: Evaluating LLM Agents on Benchtop DNA Acquisition," RAND Corporation, February 3, 2026, https://www.rand.org/pubs/research_reports/RRA4591-1.html.
92. Jasper Götting et al., "Virology Capabilities Test (VCT): A Multimodal Virology Q&A Benchmark," arXiv:2504.16137, April 21, 2025, <https://arxiv.org/abs/2504.16137>.
93. Bill Drexel and Caleb Withers, "AI and the Evolution of Biological National Security Risks: Capabilities, Thresholds, and Interventions," Center for a New American Security,

- August 13, 2024, <https://www.cnas.org/publications/reports/ai-and-the-evolution-of-biological-national-security-risks>.
94. *Adherence to and Compliance with Arms Control, Nonproliferation, and Disarmament Agreements and Commitments* (U.S. Department of State, April 2025), <https://www.state.gov/wp-content/uploads/2025/04/2025-Arms-Control-Treaty-Compliance-Report-1.pdf>; *Annual Report to Congress: Military and Security Developments Involving the People's Republic of China, 2025* (U.S. Department of Defense, December 18, 2025), <https://media.defense.gov/2025/Dec/23/2003849070/-1/-1/1/ANNUAL-REPORT-TO-CONGRESS-MILITARY-AND-SECURITY-DEVELOPMENTS-INVOLVING-THE-PEOPLES-REPUBLIC-OF-CHINA-2025.PDF>; *Military and Security Developments Involving the People's Republic of China, 2024* (U.S. Department of Defense, December 18, 2024), <https://media.defense.gov/2024/Dec/18/2003615520/-1/-1/0/MILITARY-AND-SECURITY-DEVELOPMENTS-INVOLVING-THE-PEOPLES-REPUBLIC-OF-CHINA-2024.PDF>; *Military and Security Developments Involving the People's Republic of China, 2023* (U.S. Department of Defense, 2023), <https://media.defense.gov/2023/Oct/19/2003323409/-1/-1/1/2023-MILITARY-AND-SECURITY-DEVELOPMENTS-INVOLVING-THE-PEOPLES-REPUBLIC-OF-CHINA.PDF>.
 95. *Annual Threat Assessment of the U.S. Intelligence Community* (Office of the Director of National Intelligence, February 6, 2023), <https://www.odni.gov/files/ODNI/documents/assessments/ATA-2023-Unclassified-Report.pdf>; *Updated Assessment on COVID-19 Origins* (Office of the Director of National Intelligence, October 2021), <https://www.dni.gov/files/ODNI/documents/assessments/Declassified-Assessment-on-COVID-19-Origins.pdf>; Dan De Luce, "CIA Shifts Assessment on Covid Origins, Saying Lab Leak Likely Caused Outbreak," NBC News, January 25, 2025, <https://www.nbcnews.com/politics/politics-news/cia-shifts-assessment-covid-origins-saying-lab-leak-likely-caused-outb-rcna189284>.
 96. Z.ai, "GLM-5: From Vibe Coding to Agentic Engineering," GitHub, accessed April 3, 2026, <https://github.com/zai-org/GLM-5>; "Kimi K2.5: Visual Agentic Intelligence."
 97. Jackson Kaunismaa et al., "Eliciting Harmful Capabilities by Fine-Tuning on Safeguarded Outputs," arXiv:2601.13528, January 20, 2026, <https://arxiv.org/abs/2601.13528>.
 98. Frontier AI Risk Monitoring Platform, "Evaluations: Biological Risks," Frontier AI Risk Monitoring Platform, accessed April 3, 2026, <https://airiskmonitor.net/domain/bio/evaluation>.
 99. Anson Ho and Arden Berg, "Do the Biorisk Evaluations of AI Labs Actually Measure the Risk of Developing Bioweapons?," EpochAI, June 13, 2025, <https://epoch.ai/gradient-updates/do-the-biorisk-evaluations-of-ai-labs-actually-measure-the-risk-of-developing-bioweapons>.
 100. Luca Righetti, "Dual-Use AI Capabilities and the Risk of Bioterrorism: Converting Capability Evaluations to Risk Assessments," Centre for the Governance of AI, December 12, 2025, <https://www.governance.ai/research-paper/dual-use-ai-capabilities-and-the-risk-of-bioterrorism-converting-capability-evaluations-to-risk-assessments>; Georgia Adamson and Gregory C. Allen, *Opportunities to Strengthen U.S. Biosecurity from AI-Enabled Bioterrorism: What Policymakers Should Know* (CSIS, August 6, 2025), <https://www.csis.org/analysis/opportunities-strengthen-us-biosecurity-ai-enabled-bioterrorism-what-policymakers-should>.
 101. David Bandurski, "A Proposal for AI-Powered Censorship," China Media Project, January 14, 2025, <https://chinamediaproject.org/2025/01/14/a-proposal-for-ai-powered-censorship>.
 102. Fergus Ryan et al., *The Party's AI: How China's New AI Systems Are Reshaping Human Rights* (Australian Strategic Policy Institute, December 1, 2025), <https://www.aspi.org.au/report/the-partys-ai-how-chinas-new-ai-systems-are-reshaping-human-rights>.
 103. Alex Colville, "China's AI Content Dragnet," China Media Project, March 24, 2025, <https://chinamediaproject.org/2025/03/24/chinas-ai-content-dragnet>.
 104. Chin and Lin, *Surveillance State*; Pei, *The Sentinel State*.
 105. Wenjing Zhang et al., "CHiSafetyBench: A Chinese Hierarchical Safety Benchmark for Large Language Models," arXiv:2406.10311, June 14, 2024, <https://arxiv.org/abs/2406.10311>; Hengxiang Zhang et al., "ChineseSafe: A Chinese Benchmark for Evaluating Safety in Large Language Models," arXiv:2410.18491, April 13, 2025, <https://arxiv.org/abs/2410.18491>.
 106. Shuai Bai et al., "Qwen2.5-VL Technical Report," arXiv:2502.13923, February 19, 2025, <https://arxiv.org/abs/2502.13923>; Kimi Team, "Kimi-VL Technical Report," arXiv:2504.07491, June 23, 2025, <https://arxiv.org/abs/2504.07491>; Xiaokang Chen et al., "Janus-Pro: Unified Multimodal Understanding and Generation with Data and Model Scaling," arXiv:2501.17811, January 29, 2025, <https://arxiv.org/abs/2501.17811>; "GLM-4.6V," ZhipuAI, accessed March 31, 2026, <https://docs.z.ai/guides/vlm/glm-4.6v>; *State of AI: China, Q2 2025 Highlights Report* (Artificial Analysis, 2025), <https://artificialanalysis.ai/downloads/china-report/2025/Artificial-Analysis-State-of-AI-China-Q2-2025-Highlights.pdf>; *State of AI: Q3 2025 Highlights Report* (Artificial Analysis, 2025), <https://artificialanalysis.ai/downloads/state-of-ai/2025/Q3-2025-Artificial-Analysis-State-of-AI-Highlights-Report.pdf>.
 107. Fatima Tlis, "China Uses DeepSeek, Other AI Models, for Surveillance and Information Attacks on US," Voice of America, March 4, 2025, <https://www.voanews.com/a/china-uses-deepseek-ai-for-surveillance-and-information-attacks-on-us/7996271.html>; "SenseTime Integrates DeepSeek Enterprise into SenseCore and Launches the Open-Source LazyLLM Framework," SenseTime, February 26, 2025, <https://www.sensetime.com/en/news-detail/51169384?categoryId=1072>; Emiko Matsui, "Huawei and iFlytek Unveiled Spark All-in-One Model with DeepSeek Support," Huawei Central Newsroom, March 3, 2025, <https://www.huaweicentral.com/huawei-and-iflytek-unveiled-spark-all-in-one-model-with-deepseek-support>.
 108. Lind, *Autocracy 2.0*.

109. Russell E. McGuire et al., “Exploring Artificial Intelligence-Enhanced Cyber and Information Operations Integration,” *Military Review* (March–April 2025): 8–19, <https://www.armyupress.army.mil/Portals/7/military-review/Archives/English/MA-25/AI-Cyber-Information-Operations-Integration/AI-Cyber-Information-Operations-Integration-UA.pdf>.
110. swyx (Shawn Wang) and Alessio Fanelli, “[AINews] Qwen Image 2 and Seedance 2,” Latent Space, February 11, 2026, <https://www.latent.space/p/ainews-qwen-image-2-and-seedance>.
111. QwenTeam, “Qwen-Image-2.0: Professional Infographics, Exquisite Photorealism,” Qwen Blog, February 10, 2026, <https://qwen.ai/blog?id=qwen-image-2.0>; Nick Corvino, “Seedance, Kling and the Chinese AI Video Ecosystem: A Regulatory Puzzle,” ChinaTalk, February 13, 2026, <https://www.chinatalk.media/p/seedance-kling-and-the-chinese-ai>.
112. *Disrupting Malicious Uses of Our Models: An Update, February 2026* (OpenAI, 2026), <https://cdn.openai.com/pdf/df438d70-e3fe-4a6c-a403-ff632def8f79/disrupting-malicious-uses-of-ai.pdf>.
113. *Green Cicada Network: Emerging X (Twitter) Inauthentic Account Network Powered by Generative AI* (CyberCX Intelligence, August 13, 2024), <https://connect.cybercx.com.au/Intelligence-Update-CCX-IU-2024-004>.
114. Brett Goldstein and Brett V. Benson, “The GoLaxy Documents: Inside One Chinese Company’s AI-Driven Influence Machine,” Wicked Problems Lab, Vanderbilt University, Institute of National Security, accessed March 31, 2026, <https://www.vanderbilt.edu/national-security/wicked-problems-lab/golaxy>.
115. Julian E. Barnes, “China Turns to A.I. in Information Warfare,” *The New York Times*, August 6, 2025, <https://www.nytimes.com/2025/08/06/us/politics/china-artificial-intelligence-information-warfare.html>.
116. “AI and the Future of Democratic Defense,” OpenAI Global Affairs, February 25, 2026, <https://openaiglobalaffairs.substack.com/p/ai-and-the-future-of-democratic-defense>.
117. *Disrupting Malicious Uses of AI: June 2025* (OpenAI, June 2025), <https://cdn.openai.com/threat-intelligence-reports/5f73af09-a3a3-4a55-992e-069237681620/disrupting-malicious-uses-of-ai-june-2025.pdf>.
118. Clint Watts, “China Tests US Voter Fault Lines and Ramps AI Content to Boost Its Geopolitical Interests,” Microsoft Threat Analysis Center, April 4, 2024, <https://blogs.microsoft.com/on-the-issues/2024/04/04/china-ai-influence-elections-mtac-cybersecurity>.
119. Leah Siskind and Marina Chernin, “Deepfakes on the Front Lines: Iran’s AI Disinformation Campaign,” Foundation for Defense of Democracies, March 19, 2026, <https://www.fdd.org/analysis/2026/03/19/deepfakes-on-the-front-lines-irans-ai-disinformation-campaign>.
120. PeiHsuan Huang et al., “Analysis of LLM Bias (Chinese Propaganda & Anti-US Sentiment) in DeepSeek-R1 vs. ChatGPT o3-mini-high,” arXiv:2506.01814, June 2, 2025, <https://arxiv.org/abs/2506.01814>; Xulang Zhang et al., “A Systematic Analysis of Biases in Large Language Models,” arXiv:2512.15792, March 4, 2026, <https://arxiv.org/abs/2512.15792>; Jennifer Pan and Xu Xu, “Political Censorship in Large Language Models Originating from China,” *PNAS Nexus* 5, no. 2 (2026): pgag013, <https://academic.oup.com/pnasnexus/article/5/2/pgag013/8487339>.
121. *Evaluation of DeepSeek AI Models*; “CAISI Evaluation of Kimi K2 Thinking.”
122. *Evaluation of DeepSeek AI Models*.
123. “CAISI Evaluation of Kimi K2 Thinking.”
124. Ryan et al., *The Party’s AI*.
125. Stefan Stein, “CrowdStrike Research: Security Flaws in DeepSeek-Generated Code Linked to Political Triggers,” *CrowdStrike* (blog), November 20, 2025, <https://www.crowdstrike.com/en-us/blog/crowdstrike-researchers-identify-hidden-vulnerabilities-ai-coded-software>.
126. “Translation: Interim Measures for the Management of Generative Artificial Intelligence Services.”
127. Pan and Xu, “Political Censorship in Large Language Models Originating from China”; Alex Colville, “Alibaba’s AI Bias Problem,” China Media Project, October 3, 2025, <https://chinamediaproject.org/2025/10/03/alibas-ai-bias-problem>.
128. Jan Betley et al., “Weird Generalization and Inductive Backdoors: New Ways to Corrupt LLMs,” arXiv:2512.09742, December 10, 2025, <https://arxiv.org/abs/2512.09742>.
129. Alex Colville, “Tokens of AI Bias,” China Media Project, February 9, 2026, <https://chinamediaproject.org/2026/02/09/tokens-of-ai-bias>.
130. “Chinese Artificial Intelligence Distorts Perceptions,” in *International Security and Estonia 2026* (Estonian Foreign Intelligence Service) 72–75, <https://raport.valisluureamet.ee/2026/en/6-asia/6-3-chinese-artificial-intelligence-distorts-perceptions>.
131. Gabriel Wagner et al., “State of AI Safety in China,” Concordia AI, July 2025, <https://concordia-ai.com/wp-content/uploads/2025/07/State-of-AI-Safety-in-China-2025.pdf>.
132. Alexander Wan et al., *The 2025 Foundation Model Transparency Index* (Center for Research on Foundation Models, Stanford University, December 2025), <https://crfm.stanford.edu/fmti/December-2025/paper.pdf>.
133. Jeffrey Ding, “ChinAI #351: CAICT Launches 2026 AI Safety Evaluations,” ChinAI Newsletter, March 16, 2026, <https://chinai.substack.com/p/chinai-351-caict-launches-2026-ai>.
134. Matt Pearl et al., “Delving into the Dangers of DeepSeek,” CSIS, February 24, 2025, <https://www.csis.org/analysis/delving-dangers-deepseek>; Emma Rafaelof et al., trans., “Translation: Data Security Law of the People’s Republic of China (Effective Sept.

- 1, 2021),” DigiChina, June 29, 2021, <https://digichina.stanford.edu/work/translation-data-security-law-of-the-peoples-republic-of-china>; Rogier Creemers et al., trans., “Translation: Cybersecurity Law of the People’s Republic of China (Effective June 1, 2017),” DigiChina, June 29, 2018, <https://digichina.stanford.edu/work/translation-cybersecurity-law-of-the-peoples-republic-of-china-effective-june-1-2017>.
135. Théodore Christakis, “DeepSeek and the China Data Question: Direct Collection, Open Source and the Limits of Extraterritorial Enforcement,” IAPP, February 4, 2026, <https://iapp.org/news/a/deepseek-and-the-china-data-question-direct-collection-open-source-and-the-limits-of-extraterritorial-enforcement>; Charles Owen-Jackson, “How Criminals Are Compromising AI Software Supply Chains,” IBM, November 18, 2025, <https://www.ibm.com/think/insights/cyber-criminals-compromising-ai-software-supply-chains>.
 136. Ravie Lakshmanan, “Researcher Uncovers 30+ Flaws in AI Coding Tools Enabling Data Theft and RCE Attacks,” The Hacker News, December 6, 2025, <https://thehackernews.com/2025/12/researchers-uncover-30-flaws-in-ai.html>; Ziv Karliner, “New Vulnerability in GitHub Copilot and Cursor: How Hackers Can Weaponize Code Agents,” Pillar Security, March 18, 2025, <https://www.pillar.security/blog/new-vulnerability-in-github-copilot-and-cursor-how-hackers-can-weaponize-code-agents>; Sage Lazzaro, “AI Coding Tools Exploded in 2025. The First Security Exploits Show What Could Go Wrong,” *Fortune*, December 15, 2025, <https://fortune.com/2025/12/15/ai-coding-tools-security-exploit-software>; Apostol Vassilev et al., *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations* (NIST, 2025), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-2e2025.pdf>; Owen-Jackson, “How Criminals Are Compromising AI Software Supply Chains”; Pierluigi Paganini, “Rules File Backdoor: AI Code Editors Exploited for Silent Supply Chain Attacks,” Security Affairs, March 19, 2025, <https://securityaffairs.com/175593/hacking/rules-file-backdoor-ai-code-editors-silent-supply-chain-attacks.html>.
 137. You-Hao Lai, *The Authoritarian Gaze: China’s Global Data Reach and the Systemic Risks to Democracy* (Research Institute for Democracy, Society and Emerging Technology, January 20, 2026), <https://dset.tw/en/research/the-authoritarian-gaze>.
 138. *DeepSeek Unmasked: Exposing the CCP’s Latest Tool for Spying, Stealing, and Subverting U.S. Export Control Restrictions* (The Select Committee on the Strategic Competition Between the United States and the Chinese Communist Party, United States Congress, April 16, 2025), <https://chinaselectcommittee.house.gov/sites/evo-subsites/selectcommitteeontheccp.house.gov/files/evo-media-document/DeepSeek%20Final.pdf>.
 139. NIST, “CAISI Evaluation of DeepSeek AI Models Finds Shortcomings and Risks,” press release, November 20, 2025, <https://www.nist.gov/news-events/news/2025/09/caisi-evaluation-deepseek-ai-models-finds-shortcomings-and-risks>.
 140. Kevin Poirault, “DeepSeek’s Flagship AI Model Under Fire for Security Vulnerabilities,” *Infosecurity Magazine*, January 31, 2025, <https://www.infosecurity-magazine.com/news/deepseek-r1-security/>; Paul Kassianik and Amin Karbasi, “Evaluating Security Risk in DeepSeek and Other Frontier Reasoning Models,” *Cisco Blogs*, January 31, 2025, <https://blogs.cisco.com/security/evaluating-security-risk-in-deepseek-and-other-frontier-reasoning-models>; Jason Martin et al., “DeepSh*t: Exposing the Security Risks of DeepSeek-R1,” HiddenLayer, January 30, 2025, <https://www.hiddenlayer.com/research/deepsht-exposing-the-security-risks-of-deepseek-r1>; Trent Holmes and Willem Gooderham, “Exploiting DeepSeek-R1: Breaking Down Chain of Thought Security,” Trend Micro, March 4, 2025, https://www.trendmicro.com/en_us/research/25/c/exploiting-deepseek-r1.html; Milon Bhattacharya, “Is DeepSeek’s Latest Open-Source R1 Model Secure?,” Netskope, January 31, 2025, <https://www.netskope.com/blog/is-deepseeks-latest-open-source-r1-model-secure>.
 141. Asha Sharma, “DeepSeek R1 Is Now Available on Azure AI Foundry and GitHub,” *Microsoft Azure* (blog), January 29, 2025, <https://azure.microsoft.com/en-us/blog/deepseek-r1-is-now-available-on-azure-ai-foundry-and-github>.
 142. Ali Naseh et al., “R1dacted: Investigating Local Censorship in DeepSeek’s R1 Language Model,” arXiv:2505.12625v1, May 19, 2025, <https://arxiv.org/pdf/2505.12625>.
 143. “Vertex AI Release Notes,” Google Cloud, accessed March 31, 2026, <https://docs.cloud.google.com/vertex-ai/generative-ai/docs/release-notes>; “DeepSeek-V3.1 Model Now Available Fully Managed on Amazon Bedrock,” Amazon Web Services, September 18, 2025, <https://aws.amazon.com/about-aws/whats-new/2025/09/deepseek-v3-1-model-fully-managed-amazon-bedrock>; Danilo Poccia, “Qwen Models Are Now Available in Amazon Bedrock,” Amazon Web Services, September 18, 2025, <https://aws.amazon.com/blogs/aws/qwen-models-are-now-available-in-amazon-bedrock>.
 144. “Provisions on the Management of Network Product Security Vulnerabilities,” Cyberspace Administration of China, Ministry of Public Security, and Ministry of Industry and Information Technology, July 12, 2021, translated in China Law Translate, July 14, 2021, <https://www.chinalawtranslate.com/en/product-security-vulnerabilities>.
 145. *Huawei Cyber Security Evaluation Centre Oversight Board: Annual Report 2019* (UK National Cyber Security Centre, March 28, 2019), <https://www.gov.uk/government/publications/huawei-cyber-security-evaluation-centre-oversight-board-annual-report-2019>; AFP, “‘No Evidence’ of Huawei Spying, Says German IT Watchdog,” SecurityWeek, December 17, 2018, <https://www.securityweek.com/no-evidence-huawei-spying-says-german-it-watchdog/>; *Order on Final Designation of Huawei Technologies Company as a Covered Company Posing a National Security Threat*, DA 20-690 (Federal Communications Commission, June 30, 2020), <https://docs.fcc.gov/public/attachments/da-20-690a1.pdf>.
 146. Ryan Fedasiuk, “Code Red: Is the Future of American AI Being Built in Beijing?,” American Enterprise Institute, February 6, 2026, <https://www.aei.org/foreign-and-defense-policy/code-red-is-the-future-of-american-ai-being-built-in-beijing>.

147. Stack Overflow, “2025 Developer Survey: AI,” 2025, <https://survey.stackoverflow.co/2025/ai>; Irene Zhang, “What Are Chinese People Vibecoding? Domestic vs Western Agents and App Store dominance,” ChinaTalk, February 24, 2026, <https://www.china-talk.media/p/what-are-chinese-people-vibecoding>.
148. 晓曦 [Dawn], “American Programming Products Output ‘Chinese,’” 36Kr, November 2, 2025, <https://eu.36kr.com/en/p/3535638936771456>.
149. “工业和信息化部等八部门关于印发《人工智能+制造》专项行动实施意见》的通知 [Action Plan for the ‘AI Plus’ Initiative],” National Data Administration, People’s Republic of China, January 9, 2026, https://www.nda.gov.cn/sjj/zwgk/zcfb/0112/20260107214358696030895_pc.html; Kendra Schaefer and Tom Nunlist, “The AI Plus initiative – China’s Blueprint for AI Diffusion,” Trivium China, September 4, 2025, <https://triviumchina.com/research/the-ai-plus-initiative-chinas-blueprint-for-ai-diffusion>; Geopolitechs, “China Releases ‘AI Plus’ Policy.”
150. China Open Source Observatory, “The Fourth Industrial Revolution,” accessed March 31, 2026, <https://chinaopensourceobservatory.org/glossary/the-fourth-industrial-revolution>; Keith Stouffer et al., *Guide to Operational Technology (OT) Security*, NIST Special Publication 800-82, Rev. 3 (National Institute of Standards and Technology, September 2023), <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r3.pdf>.
151. “SWE-Bench,” Vals.ai, April 9, 2026, <https://www.vals.ai/benchmarks/swbench>.
152. China Daily, “SOEs Actively Deploying DeepSeek AI Models,” *China Daily*, March 3, 2025, https://regional.chinadaily.com.cn/wic/2025-03/03/c_1075991.htm; Lea Thome, “PRC Deploys DeepSeek Across Local Governments,” *China Brief* 25, no. 6 (2025), <https://jamestown.org/prc-deploys-deepseek-across-local-governments>; Jeffrey Ding, “ChinAI #348: China’s Compute Year in Review – Frenzy, Growing Pains, and Dey Milestones,” ChinAI, February 23, 2026, <https://chinai.substack.com/p/chinai-348-chinas-compute-year-in>.
153. Kimi Team, “Kimi K2.5: Visual Agentic Intelligence,” GitHub, MoonshotAI/Kimi-K2.5 repository, January 30, 2026, https://github.com/MoonshotAI/Kimi-K2.5/blob/master/tech_report.pdf.
154. Sunil Chopra, *Supply Chain Management: Strategy, Planning, and Operation*, 7th ed. (Pearson Education, 2019).
155. China Daily, “State Grid Jiaying Launches AI Dispatch System,” World Internet Conference, December 24, 2025, https://www.wicinternet.org/2025-12/24/c_1149534.htm.
156. Xi Jiang et al., “MMAD: A Comprehensive Benchmark for Multimodal Large Language Models in Industrial Anomaly Detection,” arXiv:2410.09453, February 21, 2025, <https://arxiv.org/abs/2410.09453>; Fuxiang Sun et al., “UniPCB: A Unified Framework for Open-Ended PCB Quality Inspection,” arXiv:2601.19222, January 27, 2026, <https://arxiv.org/abs/2601.19222>; Junwen Miao et al., “AgentIAD: Tool-Augmented Single-Agent for Industrial Anomaly Detection,” arXiv:2512.13671, December 15, 2025, <https://arxiv.org/abs/2512.13671>.
157. Timothy M, “Industrial Inspection Solutions,” *RoboFlow* (blog), October 20, 2025, <https://blog.roboflow.com/industrial-inspection/>.
158. Qwen Team, “Qwen3-VL Technical Report,” arXiv:2511.21631v2, November 27, 2025, <https://arxiv.org/pdf/2511.21631>; Wu Fasang, “Qwen3-VL-8B: 多模态大模型的轻量化革命, 中小企业AI落地新范式 [Qwen3-VL-8B: A Lightweight Revolution in Multimodal Large Models, a New Paradigm for AI Implementation in SMEs],” CSDN ModelEngine, October 27, 2025, <https://modelengine.csdn.net/690c52855511483559e2b0fc.html>.
159. John VerWey, *Physical AI: A Primer for Policymakers on AI-Robotics Convergence* (CSET, February 2026), <https://cset.georgetown.edu/publication/physical-ai>.
160. Pavlo Zvenyhorodskiy and Scott Singer, *Embodied AI: China’s Big Bet on Smart Robots* (Carnegie Endowment for International Peace, November 24, 2025), <https://carnegieendowment.org/research/2025/11/embodied-ai-china-smart-robots>.
161. Yicai Global, “Unitree Technology Leads the Revaluation of the Value of the Robot Track, and These Industrial Chain Companies Usher in Super Acceleration,” February 23, 2025, https://www.yicai.com/star50news/2025_02_246797215719625850883; Brenda Goh et al., “China’s AI-Powered Humanoid Robots Aim to Transform Manufacturing,” Reuters, May 13, 2025, <https://www.reuters.com/world/china/chinas-ai-powered-humanoid-robots-aim-transform-manufacturing-2025-05-13>.
162. Kyle Chan et al., “Full Stack: China’s Evolving Industrial Policy for AI,” RAND Corporation, June 26, 2025, <https://www.rand.org/pubs/perspectives/PEA4012-1.html>.
163. Alibaba Cloud Community, “BMW and Alibaba Deepen Strategic Partnership in China, Harnessing Qwen’s AI Power to Redefine Intelligent In-Car Experiences,” March 27, 2025, https://www.alibabacloud.com/blog/bmw-and-alibaba-deepen-strategic-partnership-in-china-harnessing-qwens-ai-power-to-define-intelligent-in-car-experiences_602094.
164. “HY-WorldPlay,” GitHub, Tencent-Hunyuan/HY-WorldPlay repository, accessed March 31, 2026, <https://github.com/Tencent-Hunyuan/HY-WorldPlay>; Zhou Xian et al., “Genesis: A Generative and Universal Physics Engine for Robotics and Beyond,” Genesis Embodied AI (GitHub), December 2024, <https://genesis-embodied-ai.github.io>.
165. Xinhua News Agency, “At the 20th Collective Study Session of the CCP Central Committee Politburo, Xi Jinping Stresses: Persist in Being Self-Reliant, Be Strongly Oriented Toward Applications, and Push the Orderly Development of Artificial Intelligence,” CSET, trans., April 26, 2025, <https://cset.georgetown.edu/publication/xi-politburo-collective-study-ai-2025>.
166. Geopolitechs, “China Releases ‘AI Plus’ Policy”; “两会受权发布 | 中华人民共和国国民经济和社会发展第十五个五年规划纲要 [Outline of the 15th Five-Year Plan for the National Economic

- and Social Development of the People's Republic of China],” Xinhua News Agency, People's Republic of China, March 13, 2026, <https://www.news.cn/politics/20260313/085af5de5a-4b4268aa7d87d90817df2f/c.html>.
167. Miryam Naddaf, “More Than Half of Researchers Now Use AI for Peer Review – Often Against Guidance,” *Nature* 649, no. 8096 (2026): 273–274, <https://www.nature.com/articles/d41586-025-04066-5>.
 168. Cunshi Wang et al., “StarWhisper Telescope: An AI framework for Automating End-to-End Astronomical Observations,” *Communications Engineering* 4 (2025): art. 184, <https://www.nature.com/articles/s44172-025-00520-4>.
 169. Helen Toner et al., *When AI Builds AI: Findings from a Workshop on Automation of AI R&D* (CSET, January 27, 2026), <https://cset.georgetown.edu/publication/when-ai-builds-ai>; Ben Rank et al., “PostTrainBench: Can LLM Agents Automate LLM Post-Training?” arXiv:2603.08640, March 10, 2026, <https://arxiv.org/abs/2603.08640>.
 170. “MiniMax M2.7: Early Echoes of Self-Evolution,” MiniMax, March 18, 2026, <https://www.minimax.io/news/minimax-m27-en>.
 171. OpenAI, “Early Experiments in Accelerating Science with GPT-5,” November 20, 2025, <https://openai.com/index/accelerating-science-gpt-5>; “Claude for Life Sciences,” Anthropic, October 20, 2025, <https://www.anthropic.com/news/claude-for-life-sciences>.
 172. Gregory C. Allen, “DeepSeek, Huawei, Export Controls, and the Future of the U.S.-China AI Race,” CSIS, March 7, 2025, <https://www.csis.org/analysis/deepseek-huawei-export-controls-and-future-us-china-ai-race>; Aqib Zakaria, “How Much AI Does \$1 Get You in China vs America?” ChinaTalk, February 19, 2026, <https://www.chinatalk.media/p/how-much-ai-does-1-get-you-in-china>.
 173. “Qwen3: Think Deeper, Act Faster,” Qwen, April 28, 2025, <https://qwen.ai/blog?id=qwen3>.
 174. NIST, “CAISI Evaluation of DeepSeek AI Models Finds Shortcomings and Risks”; Ryan Fedasiuk, *China's Transition to Scalable Intelligence* (American Enterprise Institute, February 18, 2026), <https://www.aei.org/wp-content/uploads/2026/02/Chinas-Transition-to-Scalable-Intelligence-Working-Paper.pdf>; Yuan Gao, “Alibaba, Tencent, and ByteDance Offer AI Red Packets to Lure Users,” Bloomberg, February 10, 2026, <https://www.bloomberg.com/news/newsletters/2026-02-10/alibaba-tencent-and-bytedance-offer-ai-red-packets-to-lure-users>; Kyle Chan, “Podcast: China's Top AI Players and Their Differing AI Strategies,” *High Capacity*, podcast, January 29, 2026, <https://www.highcapacity.org/p/podcast-chinas-top-ai-players-and>.
 175. Jeffrey Ding, “ChinAI #349: Tokens Made in China?,” ChinAI Newsletter, March 2, 2026, <https://chinai.substack.com/p/chinai-348-tokens-made-in-china>.
 176. Caiwei Chen, “What's Next for Chinese Open-Source AI,” *MIT Technology Review*, February 12, 2026, <https://www.technologyreview.com/2026/02/12/1132811/whats-next-for-chinese-open-source-ai>.
 177. Crystal Liu, “Qwen Ecosystem Expands Rapidly, Accelerating AI Adoption Across Industries,” Alibaba Cloud Community, June 30, 2025, <https://www.alibabacloud.com/blog/602330>.
 178. swyx and Alessio Fanelli, “[AINews] Qwen3.5-397B-A17B: The Smallest Open-Opus Class, Very Efficient Model,” Latent Space, February 16, 2026, <https://www.latent.space/p/ainews-qwen35-397b-a17b-the-smallest>.
 179. Adrien Laurent, “An Overview of Chinese Open-Source LLMs (Sept 2025),” Intuition Labs, September 29, 2025, <https://intuitionlabs.ai/articles/chinese-open-source-llms-2025>; Jonathan Kemper, “China Captured the Global Lead in Open-Weight AI Development During 2025, Stanford Analysis Shows,” The Decoder, January 10, 2026, <https://the-decoder.com/china-captured-the-global-lead-in-open-weight-ai-development-during-2025-stanford-analysis-shows>.
 180. “Introducing gpt-oss,” OpenAI, August 5, 2025, <https://openai.com/index/introducing-gpt-oss>.
 181. “Qwen3: Think Deeper, Act Faster”; “Qwen3-Coder: Agentic Coding in the World,” Qwen, July 22, 2025, <https://qwen.ai/blog?id=qwen3-coder>; “Qwen3.5: Towards Multimodal Agents,” Qwen, February 15, 2026, <https://qwen.ai/blog?id=qwen3.5>; An Yang et al., “Qwen3 Technical Report,” Alibaba, arXiv:2505.09388, May 2025, <https://arxiv.org/abs/2505.09388>; Shuai Bai et al., “Qwen2.5-VL Technical Report,” Alibaba, arXiv:2502.13923, February 2025, <https://arxiv.org/abs/2502.13923>; Jin Xu et al., “Qwen3-Omni Technical Report,” Alibaba, arXiv:2509.17765, September 2025, <https://arxiv.org/abs/2509.17765>; Daya Guo et al., “DeepSeek-Coder: When the Large Language Model Meets Programming – The Rise of Code Intelligence,” DeepSeek-AI, arXiv:2401.14196, January 2024, <https://arxiv.org/abs/2401.14196>; DeepSeek-AI, “DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model,” DeepSeek, arXiv:2405.04434, May 2024, <https://arxiv.org/abs/2405.04434>; DeepSeek-AI, “DeepSeek-V3 Technical Report,” DeepSeek, arXiv:2412.19437, December 2024, <https://arxiv.org/abs/2412.19437>; DeepSeek-AI, “DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning,” DeepSeek, arXiv:2501.12599, January 2025, <https://arxiv.org/abs/2501.12599>; and Xiaokang Chen et al., “Janus-Pro: Unified Multimodal Understanding and Generation with Data and Model Scaling,” DeepSeek, arXiv:2501.17811, January 2025, <https://arxiv.org/abs/2501.17811>; ByteDance Seed Team, “Seed-OSS Open-Source Models,” GitHub, August 20, 2025, <https://github.com/ByteDance-Seed/seed-oss>; ByteDance Seed et al., “Seed-Coder: Let the Code Model Curate Data for Itself,” ByteDance, arXiv:2506.03524, June 2025, <https://arxiv.org/abs/2506.03524>; “Seed2.0,” ByteDance Seed, February 14, 2026, <https://seed.bytedance.com/en/seed2>; “Seed Models,” ByteDance Seed, accessed March 31, 2026, <https://seed.bytedance.com/en/models>; Aili Chen et al., “MiniMax-M1: Scaling Test-Time Compute Efficiently with Lightning Attention,” MiniMax, June 2025, <https://arxiv.org/abs/2506.13585>; “MiniMax M2.5,” MiniMax, accessed March 31, 2026, <https://www.minimax.io/models/text>; “Kimi K2.5: Visual Agentic

- Intelligence”; Moonshot AI, “Kimi-K2-Instruct,” Hugging Face model card, <https://huggingface.co/moonshotai/Kimi-K2-Instruct>; Kimi Team, “Kimi-VL Technical Report,” Moonshot AI, arXiv:2504.07491, April 2025, <https://arxiv.org/abs/2504.07491>; Xingwu Sun et al., “Hunyuan-Large: An Open-Source MoE Model with 52 Billion Activated Parameters by Tencent,” Tencent, arXiv:2411.02265, November 2024, <https://arxiv.org/abs/2411.02265>; Tencent Hunyuan Team, “Hunyuan-A13B Technical Report,” GitHub, June 2025, https://github.com/Tencent-Hunyuan/Hunyuan-A13B/blob/main/report/Hunyuan-A13B_Technical_Report.pdf; “Hunyuan,” Tencent, accessed March 31, 2026, <https://hunyuan.tencent.com/>; Team GLM et al., “ChatGLM: A Family of Large Language Models from GLM-130B to GLM-4 All Tools,” Zhipu AI/Tsinghua University, arXiv:2406.12793, June 2024, <https://arxiv.org/abs/2406.12793>; Z.ai, “GLM-4-32B-0414,” Hugging Face, April 2025, <https://huggingface.co/zai-org/GLM-4-32B-0414>; GLM Team et al., “GLM-4.5: Agentic, Reasoning, and Coding (ARC) Foundation Models,” Zhipu AI/Tsinghua University, arXiv:2508.06471, August 2025, <https://arxiv.org/abs/2508.06471>; GLM-V Team et al., “GLM-4.5V and GLM-4.1V-Thinking: Towards Versatile Multimodal Reasoning with Scalable Reinforcement Learning,” Zhipu AI/Tsinghua University, arXiv:2507.01006, July 2025, <https://arxiv.org/abs/2507.01006>; “GLM-5,” Z.ai, accessed March 31, 2026, <https://z.ai/blog/glm-5>; “GLM-5-1,” Zhipu AI, accessed March 31, 2026, <https://bigmodel.cn/dev/howuse/model/glm-5-1>; Meta, “Llama-3.1-8B,” Hugging Face model card, July 23, 2024, <https://huggingface.co/meta-llama/Llama-3.1-8B>; Meta, “Llama-3.1-70B,” Hugging Face model card, July 23, 2024, <https://huggingface.co/meta-llama/Llama-3.1-70B>; Meta, “Llama-3.1-405B,” Hugging Face model card, July 23, 2024, <https://huggingface.co/meta-llama/Llama-3.1-405B>; Meta, “Llama-3.2-1B,” Hugging Face model card, September 25, 2024, <https://huggingface.co/meta-llama/Llama-3.2-1B>; Meta, “Llama-3.2-3B,” Hugging Face model card, September 25, 2024, <https://huggingface.co/meta-llama/Llama-3.2-3B>; Meta, “Llama-3.2-11B-Vision,” Hugging Face model card, September 25, 2024, <https://huggingface.co/meta-llama/Llama-3.2-11B-Vision>; Meta, “Llama-3.2-90B-Vision,” Hugging Face model card, September 25, 2024, <https://huggingface.co/meta-llama/Llama-3.2-90B-Vision>; Meta, “Llama-4-Scout-17B-16E-Instruct,” Hugging Face model card, April 5, 2025, <https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E-Instruct>; Meta, “Llama-4-Maverick-17B-128E-Instruct,” Hugging Face model card, April 5, 2025, <https://huggingface.co/meta-llama/Llama-4-Maverick-17B-128E-Instruct>; Meta, “The Llama 4 Herd: The Beginning of a New Era of Natively Multimodal AI Innovation,” *Meta AI* (blog), April 5, 2025, <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>; Google, “Gemma-3-1b-it,” Hugging Face model card, March 12, 2025, <https://huggingface.co/google/gemma-3-1b-it>; Google, “Gemma-3-4b-it,” Hugging Face model card, March 12, 2025, <https://huggingface.co/google/gemma-3-4b-it>; Google, “Gemma-3-12b-it,” Hugging Face model card, March 12, 2025, <https://huggingface.co/google/gemma-3-12b-it>; Google, “Gemma-3-27b-it,” Hugging Face model card, March 12, 2025, <https://huggingface.co/google/gemma-3-27b-it>; Google, “Gemma-3n-E4B-it,” Hugging Face model card, <https://huggingface.co/google/gemma-3n-E4B-it>; The Gemini Team, “Gemini 3.1 Pro: A Smarter Model for Your Most Complex Tasks,” Google, February 19, 2026, <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-1-pro/>; “Gemini 3.1 Pro,” Google DeepMind model card, February 19, 2026, <https://deepmind.google/models/model-cards/gemini-3-1-pro/>; “Gemini 3.1 Flash Lite,” API Documentation, Google AI for Developers; <https://ai.google.dev/gemini-api/docs/models#gemini-3.1-flash-lite>; OpenAI, “Gpt-oss-120b,” Hugging Face model card, August 5, 2025, <https://huggingface.co/openai/gpt-oss-120b>; OpenAI, “Gpt-oss-20b,” Hugging Face model card, August 5, 2025, <https://huggingface.co/openai/gpt-oss-20b>; “Introducing GPT-5.4,” OpenAI, March 5, 2026, <https://openai.com/index/introducing-gpt-5-4/>; “Introducing Claude Opus 4.6,” Anthropic, February 5, 2026, <https://www.anthropic.com/news/claude-opus-4-6>; “Models and Pricing,” xAI Developer Documentation, accessed March 31, 2026, <https://docs.x.ai/developers/models>.
182. Kevin Xu, “China’s Structural Advantage in Open Source AI,” Interconnected, June 25, 2025, <https://interconnect.substack.com/p/chinas-structural-advantage-in-open>; Kevin Xu, “Chinese Open Source: A Definitive History,” Interconnected, March 6, 2026, <https://interconnect.substack.com/p/chinese-open-source-a-definitive>.
 183. Grace X. Yang, “The Openness Paradox: Open-Source AI and China’s Quest for Cyber Sovereignty,” *Dialogues on Digital Society* 1, no. 3 (2025): 261–264, <https://doi.org/10.1177/29768640251376497>; Seth Hays, “The AI Kill Switch: Dangerous Chinese Open Source,” Center for European Policy Analysis, December 15, 2025, <https://cepa.org/article/the-ai-kill-switch-dangerous-chinese-open-source>; Guy Ward-Jackson et al., “Open Source: How Middle Powers Can Build Influence in the Age of AI,” Tony Blair Institute for Global Change, February 9, 2026, <https://institute.global/insights/tech-and-digitalisation/open-source-influence-age-of-ai>.
 184. Qi Xijia, “China’s Central SOEs Forge Ahead ‘AI+’ Transformation,” People’s Daily Online, February 25, 2025, <https://en.people.cn/n3/2025/0225/c90000-20281009.html>; Li Jiaying, “SOEs Actively Deploying DeepSeek AI Models,” *China Daily*, February 26, 2025, <https://www.chinadaily.com.cn/a/202502/26/WS67be6a14a310c240449d7453.html>.
 185. Muyang Chen, “Five Questions on Chinese Development Finance and the Foreign Aid Regime,” Institute on Global Conflict and Cooperation, Global Policy At a Glance blog, October 11, 2024, <https://ucigcc.org/blog/five-questions-on-chinese-development-finance-and-the-foreign-aid-regime/>; Tania Ghossein et al., “Public Procurement, Regional Integration, and the Belt and Road Initiative,” *The World Bank Research Observer* 36, no. 2 (2021): 131–163, <https://doi.org/10.1093/wbro/lkab004>; Kevin Zhu et al., “Migration to Open-Standard Interorganizational Systems: Network Effects, Switching Costs, and Path Dependency,” *MIS Quarterly* 30, Special Issue (2006): 515–539, <https://doi.org/10.2307/25148771>; “What Are the Benefits of Migrating From an Older DCS to a New Automation System?,” Automation.com, January 6, 2014, <https://blog.isa.org/benefits-migrating-older-dcs-new-industrial-automation-system>.
 186. Behrouz Ayaz, “China’s AI Push in the Persian Gulf Region,” *The Diplomat*, November 22, 2025, <https://thediplomat.com/2025/11/chinas-ai-push-in-the-gulf-region/>; “From Desert

- to Data Hub: The UAE’s Visionary AI Strategy and China’s Crucial Role,” UAE-China Chamber of Commerce, September 27, 2025, <https://uecn.org/from-desert-to-data-hub-the-uaes-visionary-ai-strategy-and-chinas-crucial-role>; Vivek Chilukuri and Ruby Scanlon, *Countering the Digital Silk Road: Saudi Arabia* (CNAS, September 16, 2025), <https://www.cnas.org/publications/reports/countering-the-digital-silk-road-saudi-arabia>; Malcolm Moore, “Saudi Aramco Chief Says DeepSeek AI Makes ‘Big Difference’ to Operations,” *Financial Times*, March 4, 2025, <https://www.ft.com/content/0d24dcf4-b53b-48e5-b49c-99606958a96d>.
187. “Advanced AI Evaluations at AISI: May Update,” UK AI Safety Institute, May 20, 2024, <https://www.aisi.gov.uk/blog/advanced-ai-evaluations-may-update>; *Frontier AI Trends Report* (UK AI Security Institute, December 2025), <https://www.aisi.gov.uk/frontier-ai-trends-report/pdf>; Kyle Brady et al., *Bridging the Digital to Physical Divide: Evaluating LLM Agents on Benchmark DNA Acquisition* (RAND Corporation, February 3, 2026), https://www.rand.org/pubs/research_reports/RRA4591-1.html.
 188. *US AISI and UK AISI Joint Pre-Deployment Test: OpenAI* (U.S. AI Safety Institute and UK AI Safety Institute, December 18, 2024), https://www.nist.gov/system/files/documents/2024/12/18/US_UK_AI%20Safety%20Institute.%20December_Publication-OpenAIoI.pdf; “U.S. AI Safety Institute Signs Agreements Regarding AI Safety Research, Testing and Evaluation With Anthropic and OpenAI,” NIST, August 29, 2024, <https://www.nist.gov/news-events/news/2024/08/us-ai-safety-institute-signs-agreements-regarding-ai-safety-research>; “FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI,” The White House, July 21, 2023, <https://bidenwhitehouse.archives.gov/briefing-room/statements-releases/2023/07/21/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-leading-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai>.
 189. House Select Committee on the Chinese Communist Party, letter to Secretary of Commerce Howard Lutnick on U.S.-PRC AI competition, May 23, 2025, <https://chinaselectcommittee.house.gov/sites/evo-subsites/selectcommitteeontheccp.house.gov/files/evo-media-document/Letter%20-%20US-PRC%20AI%20Competition.pdf>.
 190. *Evaluation of DeepSeek AI Models*; “CAISI Evaluation of Kimi K2 Thinking”; “CAISI Evaluation of DeepSeek V4 Pro,” NIST, May 1, 2026, <https://www.nist.gov/news-events/news/2026/05/caisi-evaluation-deepseek-v4-pro>.
 191. Grace Shao, “Part I: The Gala, the Suburbs, and the ‘Months Behind’ Myth in LLM Labs,” AI Proem, February 20, 2026, <https://aiproem.substack.com/p/part-i-the-gala-the-suburbs-and-the>.
 192. Kinling Lo, “OpenAI Accuses DeepSeek of ‘Free-Riding’ on American R&D,” Rest of World, February 13, 2026, <https://restofworld.org/2026/openai-deepseek-distillation-dispute-us-china>; *Adversarial Distillation* (Frontier Model Forum, February 23, 2026), <https://www.frontiermodelforum.org/issue-briefs/issue-brief-adversarial-distillation>; Google Threat Intelligence Group, “GTIG AI Threat Tracker: Distillation, Experimentation, and (Continued) Integration of AI for Adversarial Use,” *Google Cloud* (blog), February 12, 2026, <https://cloud.google.com/blog/topics/threat-intelligence/distillation-experimentation-integration-ai-adversarial-use>; “Detecting and Preventing Distillation Attacks,” Anthropic, February 23, 2026, <https://www.anthropic.com/news/detecting-and-preventing-distillation-attacks>; OpenAI, “Updated Stakes for American-Led, Democratic AI,” memorandum to the House Select Committee on Strategic Competition Between the United States and the Chinese Communist Party, February 12, 2026, <https://cdn.openai.com/pdf/045aa967-ee96-4a09-94ee-3098ddf6db2c/OpenAI-US-House-Select-Cmte-Update-%5B021226%5D.pdf>; Nathan Lambert, “How Much Does Distillation Really Matter for Chinese LLMs?,” Interconnects, February 24, 2026, <https://www.interconnects.ai/p/how-much-does-distillation-really>.
 193. *Winning the Race*; Jake Sullivan, “Remarks by National Security Advisor Jake Sullivan at the Special Competitive Studies Project Global Emerging Technologies Summit,” The White House, September 16, 2022, <https://bidenwhitehouse.archives.gov/briefing-room/speeches-remarks/2022/09/16/remarks-by-national-security-advisor-jake-sullivan-at-the-special-competitive-studies-project-global-emerging-technologies-summit/>.
 194. Office of Public Affairs, 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, <https://www.commerce.gov/news/press-releases/2025/06/statement-us-secretary-commerce-howard-lutnick-transforming-us-ai>.
 195. Janet Egan, Spencer Michaels, and Caleb Withers, *Prepared, Not Paralyzed: Managing AI Risks to Drive American Leadership* (CNAS, November 20, 2025), <https://www.cnas.org/publications/reports/prepared-not-paralyzed>.
 196. *Winning the Race*.
 197. George I. Seffers, “Automation Aids Cybersecurity for Financial Sector,” AFCEA Signal, July 1, 2015, <https://www.afcea.org/signal-media/automation-aids-cybersecurity-financial-sector>; “Financial Services ISAC,” ISAO Standards Organization, accessed March 31, 2026, <https://www.isao.org/group/fs-isac>.
 198. Donald J. Trump, “Launching the Genesis Mission,” Exec. Order No 14363, November 24, 2025, <https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission>.
 199. UK Department for Science, Innovation and Technology et al., “Efforts to Share Best Practices on AI Measurement and Evaluations Driven Forward Through the International Network for Advanced AI Measurement, Evaluation and Science,” press release, December 9, 2025, <https://www.gov.uk/government/news/efforts-to-share-best-practices-on-ai-measurement-and-evaluations-driven-forward-through-the-international-network-for-advanced-ai-measurement>.

evalua; “FACT SHEET: U.S. Department of Commerce & U.S. Department of State Launch the International Network of AI Safety Institutes at Inaugural Convening in San Francisco,” NIST, February 4, 2025, <https://www.nist.gov/news-events/news/2024/11/fact-sheet-us-department-commerce-us-department-state-launch-international>.

200. China AI Power Report Act, H.R. 6275, 119th Cong. (2025), <https://www.congress.gov/bill/119th-congress/house-bill/6275/text>.

CNAS Editorial

DIRECTOR OF STUDIES
Katherine L. Kuzminski

PUBLICATIONS & EDITORIAL DIRECTOR
Maura McCarthy

SENIOR EDITOR
Emma Swislow

ASSOCIATE EDITOR
Caroline Steel

CREATIVE DIRECTOR
Melody Cook

Cover Art & Production Notes

COVER ILLUSTRATION
Matt Needle

PRINTER
CSI Printing & Graphics
Printed on an HP Indigo Digital Press

About the Center for a New American Security

The mission of the Center for a New American Security (CNAS) is to develop strong, pragmatic, and principled national security and defense policies. Building on the expertise and experience of its staff and advisors, CNAS engages policymakers, experts, and the public with innovative, fact-based research, ideas and analysis to shape and elevate the national security debate. A key part of our mission is to inform and prepare the national security leaders of today and tomorrow.

CNAS is located in Washington, D.C., and was established in February 2007 by cofounders Kurt M. Campbell and Michèle A. Flournoy. CNAS is a 501(c)3 tax-exempt nonprofit organization. Its research is independent and nonpartisan.

©2026 Center for a New American Security

All rights reserved.



The old national security playbook no longer applies. As emerging technologies reshape the battlefield, great power rivalries intensify, and traditional frameworks evolve, the ground is no longer settled.

America should set **New Rules**. Pragmatic leaders at home and abroad can no longer afford to provide yesterday's answers to today's national security challenges.

From AI and drone warfare to global alliances and economic security, America and its allies need New Rules to compete, deter, and win in the 21st century. The Center for a New American Security develops bold, principled national security policies so that today's leaders can set the New Rules of tomorrow.

Center for a New American Security
1701 Pennsylvania Ave NW
Suite 700
Washington, D.C. 20006

[CNAS.org](https://www.cnas.org) | [@CNASdc](https://twitter.com/CNASdc)

CEO
Richard Fontaine

Executive Vice President
Paul Scharre

Senior Vice President of Development
Anna Saito Carson

Contact Us
202.457.9400
info@cnas.org



CNAS



**New
Rules**