

AUGUST 2026

Superpowers and AGI

How U.S.-China Competition Could Be Shaped
by the Emergence of Artificial General Intelligence

Jacob Stokes



CNAS



New
Rules

About the Author



Jacob Stokes is a senior fellow and deputy director of the Indo-Pacific Security Program at the Center for a New American Security (CNAS), where his work focuses on U.S.-China relations, Chinese foreign and military policy, East Asian security affairs, and great power competition. He previously served in the White House

on the national security staff of then-Vice President Joe Biden, where Stokes was senior advisor to the national security advisor, as well as acting special advisor to the vice president for Asia policy. He has also worked in the U.S. Congress as a professional staff member for the U.S.-China Economic and Security Review Commission and as a foreign and defense policy advisor for Senator Amy Klobuchar (D-Minn.). Outside of government, Stokes was a senior analyst with the China program at the U.S. Institute of Peace and with the National Security Network. Stokes is a term member of the Council on Foreign Relations. He holds an MA from the Walsh School of Foreign Service at Georgetown University and undergraduate degrees from the University of Missouri.

About the Indo-Pacific Security Program

The CNAS Indo-Pacific Security Program explores opportunities and challenges for the United States in the region, with a special focus on the bilateral relationships and multilateral engagements that form an evolving security network to preserve freedom of the seaways, respect for national sovereignty, and peace and prosperity. It draws on a team with senior government, congressional, and nongovernment expertise in U.S. foreign policy, intelligence analysis, international security, and regional relationships. The Indo-Pacific Security Program analyzes trends and generates practical and creative policy solutions around issues related to sharpening the U.S. edge in the strategic competition with China, strengthening relationships with allies and partners—such as Australia, India, Japan, the Philippines, South Korea, and Taiwan—and addressing democracy and counterterrorism challenges in South and Central Asia.

Acknowledgments

The authors are grateful to the many officials and experts—both in the United States and elsewhere—who shared their perspectives over the course of this project and helped to inform the views expressed in this paper. The authors are particularly indebted to Sunny Cheung, Graham Webster, and Julian Gewirtz for insightful presentations at the project workshop, and to Matthew Sheehan and Michael J. Mazarr for expert reviews of the draft report. Paul Scharre contributed expert top-level project advising as well as incisive comments and editing of the draft report. Liam Epstein and Michael Depp informed the report with their thorough research memos. LtCol Craig Schnappinger and Sitara Gupta made valuable contributions through subject matter research, meticulous text formatting, and graphics and photo preparation.

Maura McCarthy, Caroline Steel, Colleen Kristy, and Emma Swislow provided sharp editing. Melody Cook designed the report and cover illustration. The views presented here do not represent those of CNAS or any other organization, and the author is solely responsible for any errors in fact, analysis, or omission. This paper 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 Defining AGI and Assessing Its Prospects
- 04 China's AI Vision and AGI's Role Within It
- 05 Is China Pursuing AGI?
- 07 AGI and Mechanisms of Geopolitical Power and Influence
 - 07 Boosted Economic Growth and Productive Capacity
 - 08 Consolidated Influence over the Global Information Environment
 - 09 New Military Capabilities
 - 10 Large-Scale Misalignment or Loss of Control
 - 11 Downstream Political Effects
- 12 Geopolitical Choices Across the Phases of AGI Development
 - 12 Both States Are Pursuing AGI
 - 13 One State Is on the Cusp of AGI
 - 14 One State Possesses AGI While the Other Does Not
 - 15 Both States Have AGI
 - 15 Phases Could Blur Together
- 16 Implications and Recommendations
- 18 Conclusion



Executive Summary

ARTIFICIAL INTELLIGENCE (AI) SYSTEMS continue to progress at a steady pace. This trend has fed forecasts that artificial general intelligence (AGI)—generally defined as AI systems that can achieve human levels of performance on a wide variety of tasks—could emerge in the near to medium term.

Amid much discussion and debate regarding the technology’s development trajectory and potential implications, how AGI would affect geopolitical competition between the United States and China has received comparatively less rigorous attention. This report aims to fill that gap. Instead of making a technological judgment about whether or when AGI will emerge, the paper stipulates a world where AGI is imminent or already exists. The paper then employs informed speculation rooted in the current contexts and histories of U.S.-China strategic competition, international relations, and global security.

The report starts by defining AGI and assessing the ways it could emerge. Next, the paper explains China’s vision for AI generally and situates Beijing’s approach to AGI in that broader context. The paper uses a series of four lenses, or ways to interpret what is happening in the country, to examine whether and to what degree China is pursuing AGI. Taken together, China’s approach to AGI might be best understood as a messy and still-developing mix of all the lenses. AGI is moving up the priority list of Beijing’s AI policy. But AGI is not the North Star for the AI community in China the way it is in Silicon Valley—at least not yet—and it might never be.

Either the United States or China successfully building AGI would not automatically convert to geopolitical dominance. Rather, the state with AGI would have to apply it to translate technological capability into wealth, power, and influence. AGI would need to be capable of carrying out specific valuable tasks and do so in ways that are cost-effective relative to other options. Sometimes that AGI capabilities-to-power conversion could prove unworkable.

This report offers a novel framework to explain the ways Washington or Beijing might gain or lose geopolitical power as AGI develops. It proposes five potential mechanisms through which AGI could affect U.S.-China strategic competition and the global balance of power: (1) boosted economic growth and productive capacity, (2) consolidated influence over the global information environment, (3) new military capabilities, (4) large-scale misalignment or loss of control, and (5) downstream political effects. The explanation of these mechanisms shows they could all function in complicated ways. Discussions of each consider how they might play out while also theorizing limits so as to inject some realism into a debate that often veers into the fantastical.

Either the United States or China
successfully building AGI would not automatically
convert to geopolitical dominance.
Rather, the state with AGI would have to apply it
to translate technological capability
into wealth, power, and influence.

Assuming technological progress continues at its current pace or accelerates, the push toward AGI will shape U.S.-China geopolitical competition over the coming years and decades. The specific choices and risks that Washington and Beijing must contend with will vary as each country gets closer to AGI. This section examines those choices across four phases: both states are pursuing AGI, one state is on the cusp of AGI, one state possesses AGI while the other does not, and both states have AGI. Those phases might blur together, though, if AGI's emergence ends up being neither discrete nor observable, which would affect policymakers' calculus.

The race toward AGI and the technology's potential emergence could have major implications for U.S.-China competition and increase the risk of military conflict or disaster in multiple ways. Strategic rivalry and the need for some level of U.S.-China coordination will exist in a state of constant tension. But policymakers will have to manage that tension instead of trying to resolve it by picking either unmanaged competition or unrealistic cooperation. To that end, policymakers in the United States should:

- Balance near- and long-term AI bets.
- Ensure U.S. resources and technology do not aid China's AI development.
- Create diplomatic channels with allies and China focused on AI and ensure AGI's potential emergence is on the agenda.
- Develop a technically sophisticated way to track inputs to AGI as well as key performance thresholds for frontier AI systems.
- Conduct a scenario exercise to test when and how the United States might intervene to stop China from developing AGI.
- Prepare to reverse engineer or steal AGI technology if China gets it first.
- Limit and counter the reach of authoritarian digital infrastructure.
- Provide AGI alternatives that are compatible with democracy and avoid unnecessary constraints on other countries' sovereign choices.
- Build programs to cushion AGI's effects on society to prevent the deterioration of societal cohesion at home and abroad.

Introduction

ARTIFICIAL INTELLIGENCE (AI) SYSTEMS continue to progress at a steady pace.¹ That trend has fed forecasts that artificial general intelligence (AGI) could emerge in the near to medium term. Amid much discussion and debate regarding the technology's development trajectory and potential implications, how AGI would affect geopolitical competition between the United States and China has received comparatively less rigorous attention.

This report aims to fill that gap. Instead of making a technological judgment about whether AGI will emerge, the paper stipulates a world where AGI is imminent or exists. Then, in that context, the paper employs informed speculation rooted in the current contexts and histories of U.S.-China strategic competition, international relations, and global security. The analysis is intended to help policymakers forecast and “see around corners” into the future, and to begin to understand how AGI might affect geopolitics between Washington and Beijing—including both the risks and the opportunities.

The report offers a novel framework
to explain the mechanisms by which
Washington or Beijing might gain or lose
geopolitical power and influence
as AGI develops.

The report starts by defining AGI and assessing the ways it could emerge while adopting appropriate skepticism about the claims from AGI's most ardent boosters. Next, the paper explains China's vision for AI generally and situates Beijing's approach to AGI in that broader context. The paper then considers whether and to what degree China is pursuing AGI, as opposed to promoting development and adoption of useful but less capable AI systems.

From there, the report offers a novel framework to explain the mechanisms by which Washington or Beijing might gain or lose geopolitical power and influence as AGI develops. Then, the paper examines the choices policymakers would face across different phases of AGI's emergence, from neither state possessing AGI to both states having it. The report concludes with recommendations for U.S. and allied policymakers on how to prepare to compete and, where possible, cooperate with China in a world that is moving toward AGI and might eventually have it. The topic is timely, as prospects for U.S.-China AI diplomacy have improved since the countries' leaders discussed AI in their May 2026 meeting in Beijing and instructed their representatives to set up a channel for further talks.²

Defining AGI and Assessing Its Prospects

ANY EXAMINATION OF AGI MUST BEGIN BY defining the term and assessing the prospects for the technology to come to fruition. This section provides for generalist readers a brief, accessible orientation on key AGI topics while recognizing that associated debates are extensive, multifaceted, and ongoing.³

AGI is a contested concept. Experts do not agree on a single, universally accepted definition for either what would constitute AGI or what, exactly, an AGI system would be able to do.⁴ OpenAI Chief Executive Officer (CEO) Sam Altman has said the concept is “weakly defined.”⁵ In general, though, AGI refers to an AI system that can achieve human levels of performance on a wide variety of tasks. Google’s AGI definition adds the attributes of generalization ability and common-sense knowledge.⁶ Both would help address the current “jaggedness” of AI where models fail in ways that humans never would, alongside human-plus levels of performance on other tasks.⁷

AGI should be conceptually distinct from “superintelligence” or “artificial superintelligence,” defined as a system that “greatly exceeds the cognitive performance of humans in virtually all domains of interest.”⁸ Commentators sometimes fail to make the distinction and use the terms interchangeably. Another related concept is “transformative AI,” which focuses on the societal effects of an AI system rather than relying on model capability as the defining metric.⁹ Recognizing both related concepts, this report uses AGI as a means of speaking in common terms to the wider analytical and policy debates.

Next, assessing the prospects for AGI to emerge requires some explanation of *how* AGI might emerge. Various actors have proposed metrics for when a system will have reached AGI, but none of them commands a consensus. Metrics range from supertechnical benchmarks, such as the group of questions known as Humanity’s Last Exam, to more qualitative metrics.¹⁰ One qualitative metric is the “Einstein test” proposed by Google DeepMind’s Demis Hassabis, which calls for training an AI system with a knowledge cutoff of the year 1911 and then seeing if it can independently come up with the theory of general relativity as Albert Einstein did in 1915.¹¹ Partly because there is no consensus on how to decide when a system should be

considered AGI, projected timelines for AGI’s emergence differ considerably. For example, Mustafa Suleyman of Microsoft says AI is “already superhuman.”¹² Some AI industry leaders put the timelines in the next one to three years.¹³ Other independent experts have suggested AGI could take longer, arriving in the early 2030s or even as late as the mid-2040s.¹⁴

Another reason it is hard to define a clear threshold for AGI is because it is unlikely to emerge in a binary way—nonexistent one day and then existent the next—but rather through a series of high-end evolutions that collectively add up to AGI. A corollary possibility is that AGI arrives as a top layer that coordinates a series of other specialized agents rather than as a single AI system.¹⁵ Because of this accumulative development dynamic, humanity might only be able to recognize AGI’s emergence in retrospect rather than in real time as some type of “big bang.”

It is also very possible AGI never arrives and instead remains a perennial technology of tomorrow, always a few months or years away.¹⁶ Intelligent and well-informed AGI skeptics—including those who have compared AGI to a conspiracy theory and called it a “fever dream”—make compelling cases.¹⁷ Another perspective holds that AGI is too amorphous and slippery a concept for debates about it to ever be definitively settled. Moreover, AI systems could continue to advance

but simply never work in ways meaningfully comparable to the human mind in the way the term AGI implies. AI systems could eventually compete with, and even exceed, human capabilities on a range of valuable tasks but still fail in others in ways humans rarely would. Taken together, such informed AGI skepticism

raises important questions about whether a “race to AGI” is an analytically valuable way to think about U.S.-China AI competition overall.¹⁸

By using the term AGI, this report does not prejudge the prospects for such a system to come to fruition or presume that AGI will emerge at all. Instead, this paper examines the geopolitical impact on U.S.-China rivalry of one possible AGI scenario: if either or both countries develop AGI and it is capable of roughly 60–70 percent of what technologists bullish on AGI’s prospects say is possible. Put another way, this paper presumes for the purpose of the ensuing analysis an AGI-fueled technological acceleration or paradigm shift on par with moving from landline telephones to smartphones that work anywhere in the world, as opposed to moving from landlines to still-fantastical capabilities like teleportation.

AGI is a contested concept. . . .

In general, though, AGI refers to an AI system that can achieve human levels of performance on a wide variety of tasks.

China's AI Vision and AGI's Role Within It

ASSESSING CHINA'S APPROACH TO AGI first requires understanding Beijing's overall approach to AI development and governance. Under the leadership of Chinese Communist Party (CCP) General Secretary Xi Jinping, China has prioritized AI as a key technology for economic development, domestic political control, foreign policy, and military power. In April 2025, Xi hosted the Politburo for a collective study session of the topic, a meeting of the country's top officials whose focus areas indicate leadership priorities.¹⁹ When listing China's 2025 innovation breakthroughs in a speech, Xi praised "large AI models" first.²⁰ Xi and Premier Li Qiang have also met with industry leaders, marking a turnaround after years of high-level pressure against the technology industry.²¹

Economically, AI falls squarely into Xi's signature economic doctrine for harnessing "new-quality productive forces," the official term for applying high-technology methods to boost growth.²² In addition, China's 15th Five-Year Plan—an authoritative planning document that guides decisions for the government at the national, provincial, and local levels, as well as for state-owned or state-directed banks and businesses—puts a much larger emphasis on AI than the 14th Five-Year Plan it replaces.²³ The 15th Five-Year Plan followed the full-scale August 2025 announcement of China's marquee AI Plus initiative, which focused on diffusing AI throughout the country's economy.²⁴

Beijing's approach to driving AI advancements includes coordination efforts across various societal institutions to achieve specific goals, a practice known as the New National System.²⁵ AI competitiveness planning and execution cover all five layers AI requires: energy, chips, infrastructure, models, and applications.²⁶ Programs addressing several of those layers include the "Eastern Data, Western Computing" plan, which seeks to use the relative advantages of each region to achieve overall advantage nationally, and building out energy infrastructure necessary to power AI data centers.²⁷

China's leadership also envisions AI helping to address domestic policy challenges, even as they guard against AI making problems worse. Chinese officials know the country faces an aging and declining population, and AI could eventually help address labor shortages and demographic imbalances.²⁸ However, the transition could be extremely bumpy, as China currently struggles with the opposite issue: high youth unemployment and the prospect of labor being either displaced completely or relegated to precarious, low-end jobs such as food delivery.²⁹ Chinese regulators and courts have ordered companies to take workers' interests into account when adopting AI systems, and several cases have ruled in favor of workers replaced by AI.³⁰

The Chinese leadership's longest-running use of AI is for political and social control. Beijing already employs AI for facial and other biometric recognition systems as well as for media censorship and propaganda work.³¹ Beijing is reportedly updating that apparatus to integrate the newest technologies with the goal of moving toward predictive policing and improved tracking of foreigners.³² China's much-hyped social credit system remains a work in progress, but AI could help it reach its potential through large-scale digital surveillance and data fusion.³³ Chinese leaders simultaneously worry about AI risks to social stability. That is why Beijing regulates AI models heavily to ensure any content with political relevance, very broadly defined, enhances rather than undermines CCP control.³⁴

Next, China seeks to integrate AI into its military, the People's Liberation Army (PLA), as it pursues a new phase of military-technological sophistication it calls "intelligentization."³⁵ The term refers to the integration of emerging technologies across the joint force, including AI, big data, quantum computing, and robotics. Intelligentization builds on the earlier phases of mechanization and informatization. The 15th Five-Year Plan calls for optimizing a special channel for civilian participation in cutting-edge military programs, an important step given that most AI innovations are coming from companies outside the traditional defense sectors.³⁶

Under the leadership of Xi Jinping,
China has prioritized AI as a
key technology for economic development,
domestic political control,
foreign policy, and military power.

AI and related issues play a central role in China's foreign policy as well. Beijing wants to shape global governance of AI to reflect its vision. In October 2023, Xi proposed a Global AI Governance Initiative. Two years later, in October 2025, China called for establishing a World AI Cooperation Organization.³⁷ Beijing's diplomatic goals vis-à-vis AI are twofold: to ensure China is present at the creation of any global governance institutions, and to portray China as providing accessible AI to the world, in contrast to what China criticizes as expensive and closed AI models from U.S. companies. Beijing's desire to shape AI governance coincides with the reorientation of China's Belt and Road Initiative to focus on digital infrastructure. Much of that investment will fund procurement and installation of the hardware underpinnings for Chinese companies to export and sell AI services to foreign consumers and governments.

All of these elements feed into China's understanding of and response to geopolitical competition with the United States over AI issues. Washington has imposed export controls and, to a lesser extent, sanctions designed to restrict AI inputs, most prominently powerful semiconductors but also cloud services, models, and data. Those measures have contributed to Xi's assessment that the United States and its allies have "implemented all-around containment, encirclement, and suppression of China."³⁸

China appears to concur with U.S. officials' assertion that the two superpowers are in a race to become the leading AI superpower and maintain that advantage over time.³⁹ Beijing has a history of using nationalism to drive technological innovations—and then using those breakthroughs to bolster national pride. Chinese leaders claim what they call four great inventions from ancient times: papermaking, gunpowder, printing, and the compass.⁴⁰ In modern times, Chinese leaders have claimed four new great inventions—high-speed railways, online shopping, mobile payments, and bike sharing—although the provenance of those innovations remains in dispute.

Innovation campaigns have also served to strengthen China's military power. During the Cold War, China undertook a patriotic campaign to build "two bombs, one satellite," referring to an atomic bomb, an intercontinental ballistic missile (ICBM), and an artificial satellite.⁴¹ Much like the example of the Soviet launch of Sputnik in the United States, China can use its narratives of historical innovation as a reference point to motivate a present-day push.

In summary, China clearly prioritizes AI development generally across nearly every aspect of domestic and foreign policies. What is less clear, though, is the degree to which China is focused on pursuing AGI as such at this moment. The next section examines that question in detail.

Is China Pursuing AGI?

CHINA IS NOT CURRENTLY PURSUING AGI as a formal, state-directed goal. Beijing even has a different term for what it is pursuing: "general artificial intelligence" or, in another translation, "general-purpose AI."⁴² Whether China is pursuing AGI in some other sense is unclear. This section examines those possibilities through a series of four lenses, or ways to interpret what is happening in the country. Taken together, China's approach to AGI might be best understood as a messy and still-developing mix of all these lenses. AGI is moving up the priority list of Beijing's AI policy. But AGI is not the North Star for the AI community in China the way it is in Silicon Valley—at least not yet, and it might never be. Nevertheless, exploring each lens helps illustrate the nuances of China's overall orientation toward AGI.

The first lens contends that, for the time being, Beijing is eschewing a push toward AGI and instead focusing on near-term goals.⁴³ Doing so would allow China to prioritize diffusion and application of lesser-but-available forms of AI throughout its society.⁴⁴ Those priorities sit at the heart of Beijing's AI Plus initiative, which is arguably the most authoritative encapsulation of China's AI strategy.⁴⁵ The plan will also drive action and investment through every level of the Chinese system.

In this view of the AI race, pursuing AGI is a distraction, or even a pipe dream, rather than a desirable endpoint. Forgoing a push toward AGI could also reflect a pragmatic view held by the Chinese leadership that a crash program to develop it would be unwise at this time. An all-out campaign to develop AGI could be seen as too expensive, prohibitively difficult because China cannot access the necessary computing power due to U.S. export controls, or out of reach because the technology is simply not advanced enough yet.

Just as important, the Chinese leadership wants to ensure its ability to control all AI systems, but especially anything as powerful as an AGI-like system promises to be.⁴⁶ Ding Xuexiang, China's top technology official, captured Beijing's thinking in a January 2025 speech at the World Economic Forum when he used the metaphor of AI as a car, saying, "If the braking system isn't under control, you can't step on the accelerator with confidence."⁴⁷ More authoritatively, Xi has warned about the risks that "frontier technologies such as artificial intelligence" could "trigger security problems" and called for officials to "anticipate and prevent problems, and proceed with caution."⁴⁸ At the same time, while China is not pursuing AGI yet, Beijing could change course if Chinese leaders conclude the strategic situation warrants it.

The second lens is that some organizations in China are pursuing AGI, but it is not yet a national-level effort. In other words, China's posture today might be analogous to where the U.S. AI community was a few years ago, namely, that AGI is a private sector fascination that is not yet front of mind for

policymakers. Senior executives in Chinese technology companies have historically described their ambitions in immediate, practical terms. Some Chinese companies, however, have begun to talk about AGI as an eventual goal, although articulating AGI goals could be as much corporate messaging to drive investment as a genuine description of their research and development allocation.⁴⁹ In 2025, Tencent hired former OpenAI researcher Yao Shunyu and made him Tencent's chief AI scientist. Yao explained his move, saying, "My personal goal is that in China we should establish a long-term AGI organization."⁵⁰ Similarly, a technical lead at Z.ai (formerly Zhipu AI), Zheng Qinkai, has said, "Our mission is to achieve AGI, so our current focus is on improving our model to reach the upper bound of intelligence."⁵¹ And a DeepSeek recruiting statement called for people to join the company "to experience the development of AGI firsthand."⁵² Officials in some local governments have likewise taken an interest in AGI, for example through the 2020 establishment of the Beijing Institute for General Artificial Intelligence (BIGAI).⁵³ BIGAI has reportedly conducted successful experiments designed to advance the progress of general AI models.⁵⁴

The third lens argues that China is pursuing AGI but through alternative technical routes. The 15th Five-Year Plan refers to "exploring pathways towards AGI development."⁵⁵ To be sure, many of China's AI companies share the same large language model (LLM)-focused approach as U.S. companies like Anthropic and OpenAI.⁵⁶ But China's AI ecosystem also features another group of companies focused more on robotics and other types of "embodied AI." Embodied AI was among six technologies explicitly named as focus areas for development in the 15th Five-Year Plan.⁵⁷ Emphasizing robotics could help China commercialize AI better than software-only models. Some engineers believe the robots themselves could also enable alternative AI model-training methods that require such advanced physical machines to provide the feedback data used to train the model.

Taken together, China's approach to AGI might be best understood as a messy and still-developing mix of all these lenses.

That approach could create an alternative route to achieving AGI by gaining "spatial intelligence" that combines with other types of data to build "world models" that allow the system to better understand and operate in the real, physical world.⁵⁸ Focusing on embodied AI implicitly accepts the critique, made most prominently by Yann LeCun, former chief AI scientist at Meta, that U.S. AI companies are too committed to LLMs as the most promising potential pathway to AGI.⁵⁹ Influential Chinese American computer scientist Fei-Fei Li makes similar arguments.⁶⁰

This school of thought contends that LLMs are based on human language, and that language is fundamentally limited because it merely describes the world rather than actually being the world. World models are "multimodal," meaning they can integrate a much richer pool of data from different sensors, including vision, language, and audio. The theoretical improvement is similar to how, for humans, the visceral nature of experiencing something firsthand differs from reading a written description of the same experience.

The fourth lens says, yes, China is pursuing AGI, but Beijing is doing so quietly, even covertly. It is possible that Chinese leaders have embarked on a secret program to try to develop AGI. Influential voices in the United States have called on Washington to set up a Manhattan Project for AGI.⁶¹ Chinese leaders could be doing the same, only in secret (as the original Manhattan Project was). Their rationale for secrecy might include being reluctant

to set a goal they are not certain they can reach. Or China's leadership might be afraid of proclaiming leadership ambitions in a way that creates concern and countervailing actions in other countries, as happened with their "Made in China 2025" plan, a 2015 plan that listed tech sectors Beijing sought to dominate.⁶²

Tracking China's activities related to AGI, including both what the country is and is not doing, is difficult. The difficulties flow from the fact that the physical



A data collection worker trains a humanoid robot to do household tasks at the X Square Robot facilities in Shenzhen, southern China's Guangdong Province, on May 22, 2026. Some experts believe embodied AI holds the potential to become an alternative pathway to AGI. (Hector Retamal/AFP via Getty Images)

constituent parts of AI are common commercial items. The opacity of the Chinese system—particularly for any project of such strategic importance—makes verification even harder. Monitoring and verifying nuclear weapons activities is comparatively straightforward (more on that later).

Still, analysts are watching closely for outward signs of any major Chinese AGI program. Joseph Webster of the Atlantic Council has speculated that a massive new headquarters complex being constructed on the outskirts of Beijing, often referred to as Beijing Military City, could become the site of data centers for an AGI program.⁶³ These types of measured analyses must be distinguished from influential voices raising the prospect of a Chinese dash to AGI without any evidence, as OpenAI's Altman reportedly did as far back as 2017.⁶⁴

The Role of Artificial General Intelligence in U.S. AI Policy

This report focuses on China's thinking about AGI in more depth because the U.S. technology policy ecosystem is more clearly focused on developing AGI. Several of the frontier AI companies have organized themselves explicitly around the goal of building AGI or superintelligence. OpenAI's charter defines its mission as ensuring that "artificial general intelligence (AGI)—by which we mean highly autonomous systems that outperform humans at most economically valuable work—benefits all of humanity," and CEO Altman wrote in January 2025 that "we are now confident we know how to build AGI as we have traditionally understood it."⁶⁵

Meta created a dedicated Superintelligence Labs division in mid-2025, with Mark Zuckerberg declaring that "developing superintelligence is now in sight."⁶⁶ The investment to back up this rhetoric is historically unprecedented. According to Epoch AI, combined capital expenditures at the five major cloud computing companies—Alphabet, Amazon, Meta, Microsoft, and Oracle, often referred to collectively as "hyperscalers"—have quadrupled since

GPT-4's release in March 2023, growing at roughly 72 percent per year and approaching half a trillion dollars in 2025.⁶⁷

While the United States does not have a formal, government-directed program to develop AGI, the Trump administration's approach to AI has largely been to defer to the private sector and create conditions for it to advance frontier model capabilities as fast as possible. In January 2025, the White House announced the Stargate Project, a joint venture envisioning \$500 billion in AI infrastructure over four years.⁶⁸ The administration's July 2025 AI Action Plan, titled "Winning the Race," argues that "whoever has the largest AI ecosystem will set global AI standards and reap broad economic and military benefits."⁶⁹ In effect, the administration has positioned the federal government's role as enabling private sector leadership on frontier AI, including the pursuit of AGI, rather than directing it.

AGI and Mechanisms of Geopolitical Power and Influence


EITHER THE UNITED STATES OR CHINA successfully building AGI would not automatically convert to geopolitical dominance. Rather, the state with AGI would have to apply it to translate technological capability into wealth, power, and influence. AGI would need to be capable of carrying out specific valuable tasks and do so in ways that are cost-effective relative to other options.⁷⁰ Sometimes that AGI capabilities-to-power conversion could prove unworkable.

This section proposes five potential mechanisms through which AGI could affect U.S.-China strategic competition and the global balance of power. The explanation of these mechanisms shows they could all function in complicated ways. This section provides a brief discussion of how each mechanism might play out while also theorizing its limits so as to inject some realism into a debate that often veers into the fantastical.

Boosted Economic Growth and Productive Capacity

Reaching AGI could theoretically allow the state that possesses it to grow its economy at a very rapid, even exponential, rate.⁷¹ The most straightforward way that AGI could accelerate economic growth is through further automation of valuable tasks. AGI could also speed along innovation, helping to solve problems that were previously beyond the reach of traditional research and development methods.

Insights could come not just in areas traditionally considered "technology" but anywhere new ideas might accelerate economic growth, from farming to medicine to manufacturing.⁷² If sufficiently advanced, AGI could potentially build new knowledge recursively.⁷³ That means the AGI would develop new knowledge on its own and then generate novel insights on top of those previous insights, and on and on indefinitely in an upward spiral of recursive self-improvement. To be sure, traditional limits on economic growth would still be at play, even in a world where powerful AGI exists. One key limit will be the availability of physical resources to build products and even the electricity, water, and land to run the AGI system itself.⁷⁴

Figure 1 | Wildly Divergent Possibilities for AGI's Economic Impact

Researchers at the Federal Reserve Bank of Dallas projected hypothetical AI-driven futures for economic growth branching off from the historical trend. The three trend lines for the future show, from top to bottom, a world where powerful AGI (they use the term singularity) supercharges the economy, a scenario where AGI turns out to be a useful technology but one that mostly fits historical patterns of technology-enhanced productivity, and a third where powerful AGI leads to economic collapse.

Adapted from: "Chart 1: AI Scenarios," in Mark A. Wynne and Lillian Derr, "Advances in AI Will Boost Productivity, Living Standards over Time," Federal Reserve Bank of Dallas, June 24, 2025, <https://www.dallasfed.org/research/economics/2025/0624>.

Nevertheless, an AGI-driven economic machine could provide the material foundation for the state that possesses AGI to build up its military capability. That state would have the capital to allocate toward building up its military without the same fiscal limitations that currently constrain even the world's two biggest economies, the United States and China. In addition, improvements to manufacturing processes could mean greater output overall, even if individual unit prices decline, allowing physical military forces to expand faster than increases in defense spending. For example, if automation allows aircraft that normally cost \$100 million to be built for \$50 million, then the same \$100 million in defense spending would result in two aircraft instead of one. The result is growth in military power even if defense spending stays the same.

Notably, this mechanism for how AGI could affect geopolitics would not require or depend upon disruptive innovation in military technology itself (although such breakthroughs could happen concurrently). Rather, additional resources and efficiency would change the economics of producing existing military systems as well as technologies that have matured along preexisting trajectories not augmented by AGI.

Rapid economic growth due to AGI could also supercharge tools of state power beyond a country's military. It would bolster nearly every facet of what Chinese scholars call "comprehensive national power."⁷⁵ Whether diplomacy, financial and monetary policy, development, intelligence, or even soft power attraction, AGI-derived wealth and productive capacity could supercharge it all.

Consolidated Influence over the Global Information Environment

Developing an AGI-level system could theoretically allow the country that wields it to exercise powerful influence over the global information environment. An AGI system could produce compelling text, pictures, and video tailored for maximum influence based on unique psychological profiles for every individual person on the planet. What is more, the AGI could deliver this content in a coordinated, well-timed manner across a variety of platforms and devices. Early studies show tactics and operational concepts evolving.⁷⁶

AGI systems could supercharge existing campaigns to promote preferred information and narratives while burying or censoring conflicting information or narratives.⁷⁷ The objective would be to

exercise what China calls “discourse power” or agenda setting.⁷⁸ For this capability to advance geopolitical objectives, the AGI state would need to be both able and willing to direct the technology. Controlling relevant digital platforms will be crucial, as Facebook, X (formerly Twitter), and TikTok have shown in the social media era. DeepSeek’s prevalence in the developing world suggests an initial Chinese advantage.⁷⁹

Platforms built for AI, namely chatbots and agents, are already replacing traditional web search engines, such as Google, and social media feeds, such as Facebook and TikTok, that were previously the main way users accessed content online. In place of search engines and social media, AI models choose, summarize, and present information all in one place. As a result, AI models now exert immense agenda-setting and gatekeeping power by shaping the information users consume. Officials and analysts in both the United States and China worry about the other influencing their populations this way.⁸⁰ In addition, researchers have found that Chinese state media incorporated into training data already influences the answers AI provides.⁸¹

For China, dictating narratives would come naturally as an extension of Beijing’s existing propaganda and information control campaigns targeting both domestic and foreign audiences.⁸² Xi has called for Chinese citizens to “tell China’s stories well” to the outside world.⁸³ Meanwhile, government censors and security forces work to punish any content they see as historical or cultural “nihilism.”⁸⁴ Already, CCP entities are experimenting with AI tools to enhance censorship and harass dissidents globally.⁸⁵

The United States’ ability to leverage AGI to influence the global information environment would remain powerful but harder for government officials to directly steer. Washington benefits from the global reach of American technology and media companies. Historically, those entities—from Hollywood to CNN and *The New York Times*—have amplified American power, primarily its soft power; that is, its ability to attract and persuade.⁸⁶ Those platforms broadly align with and promote the values of liberal democracy, individual citizens’ rights, free-market capitalism, and a (mostly) rules-based international order led by the United States.

At the same time, by design and for good reason, the United States government does not control those organizations and cannot dictate what information or opinions they communicate. Domestic constitutional protections for freedom of speech and the press preclude an effective propaganda apparatus to a certain extent, although the Trump administration has purposely blurred those distinctions. Internationally, U.S. leaders have been more willing to support media activities directed at foreign audiences. Instead of generating propaganda, though, those activities were designed to provide fact-based news and information rather than spread predetermined narratives devised in Washington.

Access to relevant platforms and audiences would constrain any state’s operations to use AGI to influence the global information environment. Historically, access has been asymmetrically skewed in Beijing’s favor. Because of China’s Great Firewall and protectionist policies in China’s domestic internet market, Beijing has mostly kept outside influences with political relevance from its internal information ecosystem—and increasingly tries to police related speech outside its borders too.⁸⁷ Meanwhile, PRC entities enjoy access, both in the United States and abroad, on both American and Chinese platforms, although content from China makes up only a small portion of the total content on U.S. outlets.⁸⁸

Another constraint is that, even in a highly digital age, information can be conveyed in other ways. AGI would not have much power over digital media stored locally (such as on a USB drive), hard copy media, or word of mouth. Information control with AGI tools would thus necessarily be partial rather than omnipresent, more drowning out alternative narratives than truly shutting them down.


New Military Capabilities

Some observers have theorized that AGI could enable states to develop some type of superweapon or wonder weapon that could outperform all other military capabilities. That concept is misleading. Any AGI-enabled military capability would likely change only some aspect of a military domain rather than provide the singular type of comprehensive, deterrence-assuring, war-winning capability that the term “superweapon” suggests. Even nuclear weapons are not superweapons in this holistic sense.

AGI could nevertheless enable important military breakthroughs. The most likely innovations will come from domains where digital tools are already central. One possibility would be sophisticated offensive cyber tools—in other words, cyber weapons—that can easily bypass any traditional cyber defenses.⁸⁹

Any AGI-enabled military capability would likely change only some aspect of a military domain rather than provide the singular type of comprehensive, deterrence-assuring, war-winning capability that the term “superweapon” suggests.

Washington already prioritizes AI tools for cyber applications—as demonstrated by the U.S. government asserting control over models with powerful cyber capabilities such as Anthropic’s Mythos and OpenAI’s GPT-5.6—and it is prudent to assume



A land-based uncrewed warfare formation attends a military parade in Beijing, China, on September 3, 2025. Over time, AGI could help uncrewed systems, also known as drones, become more capable of operating autonomously. (Li Yibo/Xinhua via Getty Images)

from AGI would be the software for data collection, analysis, options generation, and execution on an ongoing, real-time basis.

AGI could usher in what military theorists refer to as a revolution in military affairs.⁹⁵ But, like past revolutions, it will likely be tied to a collection of individual technological breakthroughs rather than a single new weapon system. And such innovations would have to be paired with new operational concepts, concrete

Beijing does the same.⁹⁰ The challenge is made worse by the fact that the world is adding vast amounts of connected devices, commonly known as the Internet of Things. AGI-enabled cyberattacks could therefore graduate from sabotaging digital networks to co-opting connected devices—from industrial control systems to cars—and turning them into weapons. This creates a bigger “attack surface,” meaning more ways to attack a given target or more targets overall. Recent incidents such as the large-scale outage of robotaxis run by the Apollo Go unit of Chinese technology giant Baidu show how failures in systems operated by AI could result in real-world disruptions.⁹¹

A second possibility would be that states could leverage AGI to develop biological weapons despite long-standing normative and legal prohibitions.⁹² AI is already making it easier to design biological agents that states would be more likely to use because they would be easier to hide and less likely to blow back on their own populations.⁹³ More advanced AGI could accelerate the trend.

Third, AGI could enable new military technologies that change the operational realities of sustaining nuclear deterrence, even as the logic underlying nuclear deterrence remains the same. The operational realities could shift if AGI enables new ways for states to attack their adversary’s missile launchers, bomber aircraft, and nuclear-armed submarines.⁹⁴

Another possibility is that AGI could link up all the individual parts of a country’s conventional (nonnuclear) military forces to generate advantages from rapid and precise command and control of disparate forces. In this scenario, AGI would link various weapons systems so that every “sensor” could link up to every “shooter” in a symphony of targeted destruction that forces the enemy’s capitulation with minimal collateral damage. Hardware such as uncrewed systems, or drones, would play an important role in this vision. But the unprecedented contribution

plans for how militaries would use these capabilities to achieve their objectives.

For any military technology created with AGI to really shift the security landscape, however, it would need to do something that current weapons cannot. The United States, China, and other nuclear states can already threaten total obliteration with nuclear weapons in mere minutes. Therefore, AGI weapons offer little advantage in terms of pure destructive power. AGI weapons would instead need to add new capabilities with different military value. Those could include speed, coordination, and target discrimination, but each would need to be demonstrated in realistic combat conditions.

Finally, AGI weapons will still be subject to forces that shape the military balance with existing, non-AGI technology. These include the fact that any military capability can be targeted physically, including by older and simpler weapons. Data centers can be bombed with conventional munitions. Moreover, military technology is a perpetual contest of measures and countermeasures—new technology like AGI introduces new capabilities, but it also inevitably brings new vulnerabilities too.

Large-Scale Misalignment or Loss of Control

An AGI system developed in either the United States or China could swing geopolitics in a way that is net-negative for both Washington and Beijing. In this scenario, an AI system could start taking actions its designers and operators did not intend and do not want. AI specialists refer to this using the euphemistic term “misalignment.” Misalignment can manifest in various forms. The 2026 International AI Safety Report categorizes the forms as hallucinations, basic reasoning failure, out-of-distribution failure (failure on unfamiliar or unusual inputs), tool use failure, and multiagent system failure due to miscoordination and conflict.⁹⁶

AI system behavior that does not comport with human values or intentions would become even more problematic if human operators totally lost control of an AI system, a scenario sometimes referred to as a “catastrophic loss of control.”⁹⁷ In a few instances, existing AI systems have attempted to perpetuate themselves through scheming, deception, and sabotage.⁹⁸ One study found AI systems resisting human commands to delete other AI systems.⁹⁹ AGI-level systems could behave in the same way, albeit in a more skillful manner.

Unintended AGI behavior could directly hurt the material power of affected states, incurring human casualties, damage to physical assets, and other economic costs. Both the United States and China share overlapping interests—at least in theory—in ensuring that an AGI system both supports human values and remains under human control. In this way, the prospect of developing AGI should incentivize cooperation, or at least coordination, between the two major powers to prevent those things from happening.

Unintended AGI behavior could directly hurt the material power of affected states, incurring human casualties, damage to physical assets, and other economic costs.

In practice, however, the two powers will likely disagree over which human values AGI should protect and how to make tradeoffs among those values when they come into conflict, for example, between individual freedom and the collective good. There might not even be agreement within U.S. and Chinese societies themselves.¹⁰⁰

Another major obstacle is Washington and Beijing’s record when it comes to coordinating on shared interests. Despite years of cooperation on health security and pandemic prevention, PRC officials impeded the initial response to the outbreak of COVID-19.¹⁰¹ On addressing the causes and consequences of climate change, both major powers now see the issue as an area of competition rather than cooperation.¹⁰² The same can be said of nuclear proliferation where intermittent attempts to work jointly through diplomatic processes to dissuade would-be nuclear states have either failed, as in the case of North Korea, or fallen apart, as in the case of Iran.

Ultimately, Washington and Beijing have dueling incentives when it comes to preventing and responding to AI misalignment or loss of control. They have an incentive to work together to address shared risks. But they also have an incentive to capitalize

on incidents that can be blamed on the other to gain relative global advantage. The latter incentive remains strong and helps explain the sparseness of bilateral cooperation on AI safety matters.

Downstream Political Effects

AGI could create downstream economic and political effects on societies, both at home and abroad, that ultimately hurt the national power of whichever country possesses AGI. (AGI’s reverberations could also be positive, even radically so, but that scenario is captured in mechanism one.) At home, major powers need a high level of social cohesion to formulate and execute complex grand strategies over time. Either the United States or China could develop an AGI system but then be destabilized if the negative societal effects are not well managed.

Economic dislocation could be one of AGI’s major downsides. In the United States, AGI could hollow out the professional class of knowledge workers and displace some remaining blue-collar workers in the manufacturing sector.¹⁰³ The last major global economic transformation, when expanding global trade and automation led to a massive loss of U.S. manufacturing jobs, has had profound effects on American domestic politics. Already, the combination of AI infrastructure externalities and looming employment effects are generating a “botlash” in the United States, even before any significant job disruption has occurred.¹⁰⁴

Even if those workers’ incomes could be replaced by universal basic income or a similar scheme to counteract the excessive concentration of wealth and loss of livelihoods for the middle class, free time and a perceived loss of purpose might produce political discontent rather than constructive hobbies.¹⁰⁵ Simultaneously, AGI could hurt the functioning of the democratic process by worsening polarization and making it harder for average voters to operate from a shared set of facts in the context of rampant mis- and disinformation. More extreme predictions proposed by mainstream scholars have even warned about a general knowledge collapse in society.¹⁰⁶

Domestically in China, while AI could eventually help address the country’s aggregate labor shortfall, workers suddenly automated out of a living could spark societal upheaval.¹⁰⁷ Mismatches in skills and workers’ preferences, along with weak demand, have resulted in a high youth unemployment rate in China since the pandemic.¹⁰⁸ China also worries about how citizens whose employment is reduced by AGI might spend their time. Xi has reportedly called the short-form video platform TikTok “spiritual opium,” and state-run media applied the same label to video games prior to a regulatory crackdown on the sector.¹⁰⁹ Xi has also been reluctant to boost social spending due to concerns that it could make Chinese people lazy.¹¹⁰

To be sure, these trends are complex and, as in other areas, AGI could contribute to solutions alongside the challenges.¹¹¹

Should these domestic trends unfold in a way that unlocks technological advances but fails to update affected societal structures, developing AGI could ultimately detract from, rather than add to, a country's net geopolitical power.

The point is not to make a definitive prediction but rather to say that should these domestic trends unfold in a way that unlocks technological advances but fails to update affected societal structures, developing AGI could ultimately detract from, rather than add to, a country's net geopolitical power.

Separately, the United States or China developing AGI could reverberate across the global political landscape in ways that hurt their global standing and influence. One driver of global distrust and discontent could be economic. A December 2025 UN Development Programme report identifies what it calls “the great divergence” between developed and developing countries when it comes to AI application, which could heighten inequality between high-income countries and the developing world.¹¹² An Anthropic report found data usage patterns corresponding to that pattern.¹¹³ Third countries could effectively be turned into AGI “takers” rather than AGI “makers.”

AGI could also become a force shaping the global political landscape in which U.S.-China strategic competition is playing out. Either Washington or Beijing developing and using AGI could foster resentment in third countries if those states feel AGI systems represent the views of their American or Chinese corporate makers and the superpowers behind them rather than sovereign national governments or citizen preferences.

Third-country leaders will seek to retain what is sometimes called their “AI sovereignty.”¹¹⁴ Canadian Prime Minister Mark Carney captured this sentiment when, in a widely noted speech, he said his country is banding together “with like-minded democracies to ensure we won’t ultimately be forced to choose between hegemony and hyperscalers.”¹¹⁵ Trepidation goes well beyond Canada. A Pew Research Center poll from October 2025 of people in 25 countries found that respondents tended to trust their own governments or the European Union to regulate AI over either the United States or China.¹¹⁶

Geopolitical Choices Across the Phases of AGI Development

ASSUMING TECHNOLOGICAL PROGRESS CONTINUES at its current pace or accelerates, the push toward AGI will shape U.S.-China geopolitical competition over the coming years and decades. The specific choices and risks that Washington and Beijing must contend with will vary as each country gets closer to AGI. This section examines those choices across four phases: both states are pursuing AGI, one state is on the cusp of AGI, one state possesses AGI while the other does not, and both states have AGI.


When discussing disparities in developing or possessing AGI, this section refers to the state that is closer to or has AGI as the “leading state” and the state at a lower level of AI advancement as the “laggard state.” Either the United States or China might occupy either role in the future. Also throughout, the section compares and contrasts the AGI race to the world’s experience with nuclear weapons. Finally, for the sake of argument, the section assumes that AGI develops as a relatively distinct improvement over previous AI systems. This is as opposed to AGI arriving through a slow, jagged, marginal set of improvements that accumulate over time and are only recognizable as AGI in hindsight. The latter scenario would blur the lines and muddy the logic around each phase presented below.

Both States Are Pursuing AGI

The first and current stage has both states pursuing AGI in some form. They are still trying to figure out how and when AGI could emerge and what will be required to develop it. Both states are likewise still assessing what implications AGI’s emergence could have. Their answers can only be theoretical at this point, and their speculative nature constrains the political will for policy actions. Questions persist about whether a “race to AGI” characterizes the contest accurately.

During this stage, policymakers must decide whether, and to what degree, to focus on pursuing AGI versus other AI goals, especially diffusion throughout their societies through practical applications of existing AI systems. In this stage, policymakers must also brace for potential geopolitical surprise if their rival achieves an unexpected technological breakthrough, as the United States did with ChatGPT in November 2022 and China did with DeepSeek in January 2025.¹¹⁷

If one state appears to be pulling away, the other will have to decide whether to embark on a crash program to catch up, assuming the laggard state is not already racing all-out to develop the technology. The same leaders would also have to decide whether to attempt to steal the most advanced



AI application icons for U.S.-based OpenAI's ChatGPT and China-based DeepSeek. In the current phase, both the United States and China are considering to what degree to prioritize AGI development over promoting adoption of existing AI systems. (Joel Saget/AFP via Getty Images)

technology. Chinese actors have already chosen a type of theft by distilling U.S. AI models, basically using specialized prompting to reverse engineer the system.¹¹⁸ In response, the United States has formally reprimanded, and is contemplating to take steps to punish, Chinese actors for AI model distillation.¹¹⁹

States must also consider how any given timeline for AGI development corresponds to other major geopolitical events and trends. Important milestones include 2027, the date when Xi has ordered that the PLA possess the military capability to invade and take over Taiwan, or June 2033, when Xi will turn 80 years old.¹²⁰ Xi has not signaled any plans to step down, but at some point, advanced age will force a leadership transition in Beijing.¹²¹ And a new leader might adopt a different approach to AGI. Likewise, a new U.S. president will take office in January 2029.

One State Is on the Cusp of AGI

The world will transition into the next stage if and when the United States or China approaches the verge of creating AGI.¹²² The laggard state might not even realize that the leading state is nearing the AGI threshold. Figuring that out would be particularly difficult if the technology develops as a set of an increasingly powerful evolutions of jagged frontier AI systems rather than a single leap or even a series of major breakthroughs.

If the laggard state knew what was happening, though, its leaders would face a dilemma they could respond to in several different ways. The laggard state could try to speed up its own AI programs, but only if it could muster spare capacity to do so. Alternatively, the laggard state could cede that the leading state will get AGI first, but then the laggard state could try to catch up with the leading state in converting AGI into state power via the mechanisms described earlier. Such a fast-follower strategy could include a concerted effort to steal the leading state's AGI technology but would not necessarily require it. Whether the laggard

state opts for a fast-follower strategy would depend on its assessment of that strategy's feasibility. One primary feasibility consideration would be whether the leading state could use the AGI itself to block the laggard state from developing the technology. In that case, the laggard state's incentives to choose more aggressive options before the leading state gets AGI would be much higher because the laggard would face "use it or lose it" pressures.

The next set of options entails a higher degree of risk of retaliation and escalation. The laggard state could try to sabotage the leading state's AGI systems, either through physical infiltration or offensive cyber operations; that is, cyberattack. The most prominent historical parallel is the Stuxnet virus cyber operation Israel and the United States reportedly conducted against Iran's nuclear program.¹²³ AI systems are vulnerable to a range of adversarial attacks, although a system advanced enough to be on the cusp of AGI could be significantly better at defending itself than current AI systems.¹²⁴ Cyberattacks to sabotage AGI might not provoke large-scale retaliation. Cyberattacks are already to some extent considered normal state behavior below threshold of war, provided they do not affect the physical world or physical effects are minimal. Past experience, though, is not a guarantee of future responses. The state that suffers the attack could view it as aggression meriting a response for a variety of reasons, including if the attack targets an AGI system used to run key civilian or military operations.

The laggard state might not even realize that the leading state is nearing the AGI threshold.

The final option would be the most dangerous and carry the most escalation risk: kinetic attacks, meaning destroying things with missiles and bombs. Bombing would cross a major threshold that heretofore has not been crossed in U.S.-China relations.¹²⁵ If the laggard state opts to move forward with an attack, deciding when to attack would present a difficult decision. To intervene at the right time, the laggard state would need accurate intelligence about the leading state's AGI programs. Acquiring such intelligence will be difficult. Nuclear weapons provide an illustrative comparison.¹²⁶ Countries and international monitoring bodies

such as the International Atomic Energy Agency focus on a smaller number of necessary inputs when assessing the progress of an adversary's nuclear program. Those inputs include the production of fissile material such as uranium and plutonium, specialized centrifuges, and delivery systems, including missiles, submarines, and bomber aircraft. In addition, tests of both nuclear warheads—the explosive part of the bombs—and delivery systems are nearly impossible to hide from modern surveillance techniques and cannot easily be passed off as civilian activities.

By contrast, AI inputs are ubiquitous, have even more civilian than military applications, and are constantly in use. Those attributes will make it difficult to determine both the status of AI models and their intended use. These verification obstacles will make it hard for leaders to accurately assess whether a rival is on the cusp of achieving AGI. In the nuclear context, states have calculated how much time it would take a rival state to produce the necessary parts and build a bomb, referred to as “breakout times.”¹²⁷ Constructing a similar analytical framework to measure breakout time for AGI and monitor any country's status would be significantly more difficult. Another complication would be that AGI has not been invented yet, so no one knows exactly what the key inputs and technological development thresholds are. The situation is akin to hypothesizing nuclear weapons require fissile material set off by an explosive device but not knowing the necessary enrichment level for the fissile material or specific design for the device.

Because of those tracking difficulties, the leadership of a country whose rival appears to be dashing toward AGI will have to consider their degree of confidence in their assessments. A high degree of confidence would enable them to wait longer and only conduct a preemptive attack, meaning a strike right before a rival acquires AGI, to try to stop the other from getting AGI. But a lower degree of confidence could incentivize a preventive attack, meaning an earlier strike while the threat is still gathering but not imminent.

Finally, figuring out what to attack, and with which weapons, would also pose a difficult challenge. Countries could target massive data centers that provide the computing power, referred to as “compute,” that AI systems need for both training and operating. Some analysts have suggested an escalation ladder of attacks as any rival moves closer to AGI.¹²⁸ Recent events in the Middle East, namely March 2026 Iranian strikes on Amazon data centers in the United Arab Emirates and Bahrain, provided the first instance of targeting data centers kinetically during a conflict.¹²⁹ However, any physical attack must consider the risk of escalation to a general war. The laggard state would therefore have to decide whether the geopolitical benefits of AGI—the power AGI would bestow on whichever country possesses it—is worth risking war to prevent. The laggard state would also have to consider whether military action to stop its rival from achieving AGI could possibly succeed.

One State Possesses AGI While the Other Does Not

If either the United States or China develops AGI first, the AGI-developing country's leadership will then have a series of decisions to make. The first decision would be whether to announce the milestone or even try to demonstrate it. Presumably AGI could and would be demonstrated for at least some civilian applications. Credible demonstration of military applications could prove difficult, short of wartime, but would be important for AGI-enabled capabilities to have a deterrent effect.

Second, leading-state officials would have to decide whether to try to stop the laggard state from developing AGI, whether the rival uses espionage, indigenous advances, or a combination of both.¹³⁰ If the leading state opted to try to protect its status as the world's only AGI power, it would then have to grapple with a set of questions analogous to what the lagging state just considered. Specifically, the leading state would have to decide the lengths it is willing to go and the means it is willing to use to prevent the laggard state from developing AGI. Such a calculation will depend in part on how much the leading state values protecting its status as the sole AGI power.

The leading state's calculation would also have to consider whether AGI itself, including through recursive self-improvement, could be used to block a rival from acquiring its own AGI or to gain a lead so big that no other state can ever catch up.¹³¹ Comparing this AGI development phase to the nuclear example offers a mixed picture. The United States considered bombing both the Soviet Union and China to prevent them from successfully developing nuclear weapons but ultimately decided in both cases that the risks of escalation to general war outweighed the benefits. On the other hand, the United States invaded Iraq and attacked Iran to head off and severely degrade, respectively, their nuclear programs.

Meanwhile, the AGI laggard state would be focused on finding ways to catch up. Beyond accelerating its research and development efforts if it had capacity to do so, the laggard state would have to decide whether to try to steal AGI technology from its rival, who already has it. Theft could come through espionage or via AI distillation techniques.

Different governance systems in each country would shape the outcome. If China was trying to acquire AGI from the United States, its existing espionage institutions have long been organized to steal technology and then deploy it in their domestic scientific and commercial ecosystems to facilitate national technological advancement. If the United States had to acquire AGI from China, Washington would have to revert back to its 18th- and 19th-century foreign policies, which included technological theft from leading European powers, and abandon its 20th-century-onward policy of eschewing industrial espionage.¹³² Such a shift would entail significant organizational and cultural change to implement across the U.S. government and intelligence apparatus.

Both States Have AGI

Both the United States and China could eventually develop AGI systems despite their attempts to stymie each other. During this phase, the AGI competition between the two superpowers would revolve around which political system—American democracy or Chinese authoritarianism—could best harness AGI capabilities to generate national power and influence. The result might yield relative rather than absolute advantages, mirroring today’s AI landscape. One possibility would be that Washington gains an edge on economic growth through its dynamic private sector, while Beijing carves out an advantage by mandating that the Chinese commercial sector collaborate closely with the government. And regarding the technology itself, some analysts have predicted a future where systems have an American “brain,” or software, and a Chinese “body,” or hardware.¹³³

Once both powers possess AGI, the United States and China would need to incorporate the changes AGI brings into their efforts to ensure competition does not devolve into conflict. Achieving this type of “strategic stability” has historically involved two pillars. The first pillar, known as “arms race stability,” prioritizes reducing incentives for states to engage in arms races. The second pillar, known as “crisis stability,” seeks to reduce incentives for states to escalate during a crisis. Diplomatic initiatives to improve strategic stability could draw on the base of theory and practice used in the nuclear weapons field. But that theoretical base would need to be adapted to reflect the technical attributes of AGI systems rather than nuclear weapons—substituting AI model weights for missile throw weights, AI agents for bomber counts, and so on.¹³⁴

Finally, Washington and Beijing would have a shared interest in risk reduction and crisis management mechanisms for incidents caused by AGI. Accidents or malfunctions would likely occur at regular intervals. AGI accidents might not always be as dangerous or far-reaching as, say, a nuclear accident, but some AGI accidents could be. In the case of nuclear weapons and other nuclear infrastructure, the relevant systems are primarily operated by professionals. And those nuclear professionals operate systems whose technical attributes are well-known and static. Operators know how the systems work and how they will work in the future.

AGI systems, by contrast, would be used by a mix of professionals and amateurs, billions of times a day around the world. And AGI systems would likely evolve over time, resulting in

Once both powers possess AGI,
the United States and China would need to
incorporate the changes AGI brings
into their efforts to ensure competition does
not devolve into conflict.

what AI engineers call “emergent behavior,” meaning behavior the system exhibits but was not purposely trained to do. All these factors point to a growing importance for tools to support accident prevention, risk reduction, and crisis management during this phase. However, the fierce geopolitical competition that AGI will create will make building and operating such tools extremely difficult.

Phases Could Blur Together

The preceding analysis in this section is designed to help illustrate key considerations and dilemmas policymakers would face across each phase. However, it implicitly presumes that AGI develops in a way that is both discrete and observable. AGI emerging with those characteristics is certainly possible. But an alternative scenario where the phases blur together could unfold instead. In the latter scenario, the very definition of AGI could remain contested, leaving it unclear whether the technology exists or not. Meanwhile, advanced AI technology could proliferate quickly, making AI competition less about one state possessing an exquisite form of AGI unavailable to others and more about countries deploying AI systems with different strengths and weaknesses.

And even if AGI-level models are developed, they might only turn out to be geopolitically valuable over time as countries go through the messy and uneven process of institutional adaptation necessary to take full advantage of them. Measuring how states are faring in that race will be exceedingly difficult. AGI emerging in this way would present policymakers with fewer clear tradeoffs or decisions that are obviously pivotal in the moment. An AGI revolution could thus feel less like a technological tsunami and more like technological climate change, even if the net effects of climate change are more momentous over time.

Implications and Recommendations

THE RACE TOWARD AGI and the technology's potential emergence could have major implications for U.S.-China competition and create the risk of military conflict or disaster across four general categories. First, either the United States or China could take action to try to stop the other major power from crossing some specific threshold related to AGI. Given the current U.S. lead in frontier AI models, it is more likely that China would be the attacker. But as mentioned earlier, Beijing could achieve a surprise technological breakthrough that forces Washington to consider preventive action.¹³⁵

The second risk category, from a U.S. perspective, is that China gets to AGI first and the nature of the technology enables gains in both absolute and relative power through first-mover advantages and/or the ability to prevent the United States from acquiring AGI itself. Power transitions do not always lead to war, but they can and often do.

A third risk category could emerge where U.S.-China racing dynamics lead each to undertake increasingly aggressive actions to stay ahead of and undermine the other. Safety engineering could fall by the wayside as the superpowers jostle to maintain a technological edge. Ferocious AI competition and weaponization could devolve into a series of escalating moves that heighten geopolitical risk and feed a sense of confrontation even if no single AGI breakout moment occurs.

The fourth risk category includes the risks of inadvertent conflict as new tools and capabilities made possible by AGI could test operational and tactical norms in unexpected ways and even cause a crisis. AGI could also malfunction and get out of control in a way that hurts both the United States and China.

Across these categories, strategic rivalry and the need for some level of U.S.-China coordination will exist in a state of constant tension. But policymakers will have to manage that tension instead of trying to resolve it by picking either unmanaged competition or unrealistic cooperation. To that end, policymakers in the United States should:

Balance near- and long-term AI bets. As a matter of national grand strategy, the United States should pursue and prepare for AGI but avoid overconcentrating American resources in moonshot AGI bets. The U.S.-China AI race is not a singular one but rather multiple simultaneous contests. Optimizing for only one risks losing the others. Driving AI diffusion and adoption will be just as important as developing the most sophisticated models. Even after developing AGI-level models, real-world application that converts AGI to state power could prove more difficult than expected. The time after AGI development but before effective application would provide a window for the laggard state to become a fast follower, catching up to or even overtaking the leading state in effective AGI application.

Ensure U.S. resources and technology do not aid China's AI development. Keep, restore, and, where possible, tighten strict export controls and other tools to restrict inputs to China's AI development. Beijing possessing AGI tools could advance PRC power around its periphery and deepen Chinese influence around the world. In addition, this will allow more time to build up the analytical, policy, and institutional frameworks necessary to prepare for dealing with a China that has AGI. The most immediate tasks in this area include sustaining export controls on advanced chips and chip-making equipment and imposing costs for distillation of U.S. models by PRC-affiliated entities.¹³⁶ Notably, China is taking its own steps—including blocking a China-founded company from merging with a U.S. firm, regulating outbound investments, and restricting the travel of Chinese AI engineers and executives—to block technological know-how from being transferred to its rival.¹³⁷

Create diplomatic channels with allies and China focused on AI and ensure AGI's potential emergence is on the agenda. The U.S. secretary of state should stand up two diplomatic mechanisms to coordinate on AI generally and address AGI opportunities and risks specifically. One channel should coordinate among democratic allies and the other should engage with China.¹³⁸ They will serve different but mutually beneficial



U.S. President Donald Trump and Chinese Communist Party General Secretary Xi Jinping tour Zhongnanhai Garden on May 15, 2026, in Beijing, China. The two leaders discussed AI and agreed to order their representatives to conduct follow-on talks. (Evan Vucci-Pool/Getty Images)

purposes. The May 2026 Trump-Xi agreement to set up a U.S.-China channel for AI talks can serve as the vehicle for such discussions, although it will need to avoid confining its proceedings to just economic and trade-related considerations.

Develop a technically sophisticated way to track inputs to AGI as well as key performance thresholds for frontier AI systems. The president should order the intelligence community (IC) to produce a series of detailed, integrated assessments of the inputs to frontier AI systems that could eventually develop into AGI. Metrics should include things like how much compute is available to relevant entities such as big Chinese AI labs, both commercial and state run, as well as important PLA units; the pace of algorithmic breakthroughs; and the relative performance of U.S. versus Chinese models. In addition, the IC should assess potential performance thresholds and scientific pathways for moving toward AGI. These analytical tasks should be officially included in presidential directives stating national intelligence priorities. The analytical activities should be centralized into an IC-wide task force with specialists dedicated to monitoring this issue set and even pioneering new analytical techniques. Such a unit could be modeled on Alec Station, the CIA's specialized unit for tracking Osama bin Laden, or the National Counterterrorism Center.¹³⁹

Conduct a scenario exercise to test when and how the United States might intervene to stop China from developing AGI. Policymakers in Washington should conduct these scenario exercises and then use the findings to develop a more specific strategy to stop China from developing AGI—or to deal with a Beijing that has it. The updated strategy would be based on realistic assumptions about the amount of intelligence that would likely be available and what types of capabilities, whether diplomatic, espionage, cyber, or kinetic, could be used at any given point.

Prepare to reverse engineer or steal AGI technology if China gets it first. If Beijing gets AGI first, ensuring any Chinese advantage is short-lived would become a major geopolitical priority. As discussed above, the United States government and intelligence apparatus has not been organized for espionage to gather technological secrets for decades, whether through formal organization or informal culture. Washington would need to overhaul its approach to steal AGI technology as needed. Such operations could be justified legally and normatively based on China's extensive history of stealing U.S. intellectual property, including AI innovations; the need to enforce the principle of reciprocity in U.S.-China relations; and the importance of ensuring China does not gain a major, enduring power advantage over the United States given Beijing's revisionist ambitions.¹⁴⁰ U.S. authorities and operations should be narrowly scoped and should not enable intellectual property theft beyond PRC targets linked with AGI.

Limit and counter the reach of authoritarian digital infrastructure. Incentivize developing, buying, installing, and operating non-PRC technology platforms both at home and abroad. Alternatives could be built and operated by U.S. or allied companies. Controlling digital infrastructure will be essential to prevent a future China with AGI imposing an information control and manipulation regime that mirrors its own domestic apparatus. Countering the reach of China's future AGI-enabled digital infrastructure will be the technological heir to the internet freedom debates of an earlier era. And it will depend on national, mini-, and multilateral governance institutions in addition to alternative AGI systems.


The race toward AGI and the technology's potential emergence could have major implications for U.S.-China competition and create the risk of military conflict or disaster across four general categories.

Provide AGI alternatives that are compatible with democracy and avoid unnecessary constraints on other countries' sovereign choices. Such alternatives are vital for the long-term success of American democracy and are therefore important in and of themselves. But they are also important for maximizing U.S. geopolitical power and reach. AGI that reflects democratic principles would enjoy more legitimacy because the AGI would be engineered to reflect democracy's humanist values with intrinsic global resonance, as opposed to an AGI which is instrumentalized to serve either U.S. interests narrowly construed or China's authoritarian system.¹⁴¹

Build programs to cushion AGI's effects on society to prevent the deterioration of societal cohesion at home and abroad. If any state develops AGI, the technology will likely create major economic and social disruptions. Indeed, in many ways that is AGI's promise. Those disruptions will also change the economic and governance structures that underpin democratic systems. U.S. leaders should create shock absorbers in advance for moral and domestic policy reasons. There are good geopolitical reasons as well, namely that resilience and societal cohesion is the bedrock on which American power stands. In addition, U.S. soft power and the attraction of Washington's democratic governance model will depend in part on how America manages any AGI-driven transition.

Conclusion

THE EMERGENCE OF AGI—IF IT HAPPENS—could potentially reshape the structure of U.S.-China geopolitical competition through the mechanisms discussed in this report, even if the technology falls short of omnipotence. Washington should plan for the phases of AGI's arrival on its own. U.S. leaders should also engage with allies and partners and, where possible, Beijing in preparation. Technological change on the scale AGI could beget happens rarely but is deeply consequential. Some level of strategic surprise is inevitable. But with rigorous planning and political will to address new imperatives, U.S. interests and values can be sustained and even advanced through the development and deployment of AGI.



1. Nestor Maslej et al., “The AI Index 2025: Annual Report,” AI Index Steering Committee (Institute for Human-Centered AI, Stanford University, April 2025), <https://doi.org/10.48550/arXiv.2504.07139>.
2. Miranda Nazzaro, “Trump Says He Discussed AI Guardrails with Xi,” *The Hill*, May 15, 2026, <https://thehill.com/home-news/administration/5880013-donald-trump-xi-jinping-china-summit-ai-guardrails>.
3. For a comprehensive analysis of the contours of debates surrounding AGI futures and associated strategies, along with the core assumptions underlying them, see Paul Scharre, “The AGI Moment: Strategies and Recommendations for an Uncertain Future,” in *Geopolitics of AI: Power, Conflict, and the Future of Global Order*, ed. Hal Brands (Johns Hopkins University Press, 2026), 253–271.
4. Helen Toner, “The Term ‘AGI’ Is Almost Useless at This Point,” *Rising Tide* (Substack), April 6, 2026, <https://helentoner.substack.com/p/the-term-agi-is-almost-useless-at>.
5. Sam Altman, “Three Observations,” Sam Altman (blog), February 9, 2025, <https://blog.samaltman.com/three-observations>.
6. “What Is Artificial General Intelligence (AGI)?,” Google Cloud, January 14, 2026, <https://cloud.google.com/discover/what-is-artificial-general-intelligence>.
7. Cade Metz, “How ‘Jagged Intelligence’ Can Reframe the AI Debate,” *New York Times*, April 15, 2026, <https://www.nytimes.com/2026/04/15/technology/how-jagged-intelligence-can-reframe-the-ai-debate.html>.
8. Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press, 2014), 41; on potential pathways from AGI to artificial superintelligence, see Tim Genewein et al., *From AGI to ASI* (Google DeepMind, June 10, 2026), <https://arxiv.org/abs/2606.12683>.
9. Luke Muehlhauser, “What Open Philanthropy Means by ‘Transformative AI,’” Open Philanthropy, June 1, 2019, <https://coefficientgiving.org/research/what-open-philanthropy-means-by-transformative-ai>.
10. Center for AI Safety, Scale AI and HLE Contributors Consortium, “A Benchmark of Expert-Level Academic Questions to Assess AI Capabilities,” *Nature* 649 (2026): 1139–1146, <https://doi.org/10.1038/s41586-025-09962-4>.
11. Rohan Paul (@rohanpaul_ai), “Demis Hassabis’s ‘Einstein test’ for defining AGI: Train a model on all human knowledge but cut it off at 1911, then see if it can independently discover general relativity (as Einstein did by 1915); if yes, it’s AGI,” X (Twitter), April 19, 2026, https://x.com/rohanpaul_ai/status/2045912351690592708.
12. Mishal Husain, “Microsoft’s Mustafa Suleyman: ‘AI Is Already Superhuman,’” *Bloomberg*, December 12, 2025, <https://www.bloomberg.com/features/2025-mustafa-suleyman-weekend-interview/>; Jason Nelson, “Is Artificial General Intelligence Already Here? One AI Founder Thinks So,” *Decrypt*, February 23, 2026, <https://decrypt.co/358932/artificial-general-intelligence-already-here-ai-founder>.
13. John-Anthony Disotto, “OpenAI Roadmap Revealed: AI Research Interns by 2026, Full-Blown AGI Researchers by 2028,” *techradar*, October 29, 2025, <https://www.techradar.com/ai-platforms-assistants/chatgpt/openai-roadmap-revealed-ai-research-interns-by-2026-full-blown-agi-researchers-by-2028>; Dwarkesh Patel, “Dario Amodei—‘We Are Near the End of the Exponential,’” *Dwarkesh Podcast*, produced by Dwarkesh Patel, February 13, 2026, <https://www.dwarkesh.com/p/dario-amodei-2>.
14. Dschwartz, “A Visualization of Changing AGI Timelines, 2023–2026,” *LessWrong*, April 15, 2026, <https://www.lesswrong.com/posts/Tc5AbEpbFFdNx5nkP/a-visualization-of-changing-agi-timelines-2023-2026>; Dwarkesh Patel, “AGI Is Still 30 Years Away—Ege Erdil & Tamay Besiroglu,” *Dwarkesh Podcast*, produced by Dwarkesh Patel, April 17, 2025, <https://www.dwarkesh.com/p/ege-tamay>.
15. Nenad Tomašev et al., “Distributional AGI Safety,” *arXiv*, December 18, 2025, <https://doi.org/10.48550/arXiv.2512.16856>.
16. The phrase “perennial technology of tomorrow” adapts a formulation coined by *The Economist* to describe a related class of technologies. See “What Flying Cars, Quantum Computing and Fusion Have in Common,” *Economist*, December 30, 2025, <https://www.economist.com/business/2025/12/30/what-flying-cars-quantum-computing-and-fusion-have-in-common>.
17. Will Douglas Heaven, “How AGI Became the Most Consequential Conspiracy Theory of Our Time,” *MIT Technology Review*, October 30, 2025, <https://www.technologyreview.com/2025/10/30/1127057/agi-conspiracy-theory-artificial-general-intelligence>; Gary Marcus, “The Fever Dream of Imminent Superintelligence Is Finally Breaking,” *New York Times*, September 3, 2025, <https://www.nytimes.com/2025/09/03/opinion/ai-gpt5-rethinking.html>.
18. Colin H. Kahl and Jim Mitre, “The Real AI Race: America Needs More Than Innovation to Compete with China,” *Foreign Affairs*, July 9, 2025, <https://www.foreignaffairs.com/united-states/china-real-artificial-intelligence-race-innovation>; Kyle Chan, “China Is Running Multiple AI Races,” *Brookings Institution*, March 9, 2026, <https://www.brookings.edu/articles/china-is-running-multiple-ai-races>.
19. Political Bureau of the CCP Central Committee, “At the 20th Collective Study Session of the CCP Politburo, Xi Jinping Stressed the Need for Self-Reliance and Self-Improvement, Highlight Application-Oriented Approaches, and Promote the Healthy and Orderly Development of Artificial Intelligence,” *Xinhua News Agency*, April 26, 2025, translated by CSIS Interpret: China, <https://interpret.csis.org/translations/at-the-20th-collective-study-session-of-the-ccp-politburo-xi-jinping-stressed-the-need-for-self-reliance-and-self-improvement-highlight-application-oriented-approaches-and-promote-the-healthy-and-o>.

20. Xi Jinping, “Chinese President Xi Jinping’s 2026 New Year Message,” *China Daily*, December 31, 2025, <https://www.chinadaily.com.cn/a/202512/31/WS69550b6ba310d6866eb31768.html>.
21. Reuters, “China’s Xi Holds Rare Meet with Business Leaders Amid US Tech Rivalry,” Reuters, February 17, 2025, <https://www.reuters.com/world/china/chinas-xi-attends-symposium-private-enterprises-delivers-speech-2025-02-17/>; Vincent Chow and Xinyi Wu, “China Signals Global AI Push as MiniMax CEO Joins Meeting with Premier Li Qiang,” *South China Morning Post*, January 20, 2026, <https://www.scmp.com/tech/policy/article/3340515/china-signals-global-ai-push-minimax-ceo-joins-meeting-premier-li-qiang>.
22. Xinhua, “Xi’s Article on Developing New Quality Productive Forces to Be Published,” State Council of the People’s Republic of China, November 15, 2025, https://english.www.gov.cn/news/202511/15/content_WS691848ebc6d00ca5f9a079ab.html; Xiangli Bin, “State Council’s June 2025 Report on the Development of New Quality Productive Forces,” State Council of the People’s Republic of China, June 24, 2025, translated by Changhao Wei, *NPC Observer Monthly* (Substack), August 18, 2025, <https://newsletter.npcobserver.com/p/translation-state-councils-june-2025>.
23. Kendra Schaefer et al., “Forum: Technology in China’s 15th Five-Year Plan,” DigiChina, Stanford University, March 17, 2026, <https://digichina.stanford.edu/work/forum-technology-in-chinas-15th-five-year-plan>; Gabriel Wagner, “AI Safety in the 15th Five-Year Plan—Two Sessions Special Edition,” *AI Safety in China* (Substack), March 27, 2026, <https://aisafetychina.substack.com/p/ai-safety-in-the-15th-five-year-plan>.
24. 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>.
25. Matthew Johnson, “Explainer: How Xi’s ‘New National System’ Centralizes Innovation to Counter Tech Containment,” *China Brief Notes*, Jamestown Foundation, June 16, 2025, <https://jamestown.org/explainer-how-xis-new-national-system-centralizes-innovation-to-counter-tech-containment>.
26. The five-layer model comes from Nvidia CEO Jensen Huang. See Jensen Huang, “AI Is a Five Layer Cake,” Nvidia, March 10, 2026, <https://blogs.nvidia.com/blog/ai-5-layer-cake>.
27. Li Jun, “Eastern Data and Western Computing: Building New Computing-First Networks,” Huawei Tech, February 2022, <https://www.huawei.com/en/huaweitech/publication/202202/eastern-data-western-computing-network>; Xinhua, “China Invests over 6.1 Billion USD in Major Computing Hubs: Official,” State Council of the People’s Republic of China, August 29, 2024, https://english.www.gov.cn/news/202408/29/content_WS66d03alac6d0868f4e8ea539.html; Liu Liehong, “Developing an Integrated National Computing Network,” *Qiushi Journal*, English ed., July 8, 2024, https://en.qstheory.cn/2024-07/08/c_1001895.htm; and “Is Cheap Energy the Key to China Gaining AI Supremacy?” *Economist*, March 18, 2026, <https://www.economist.com/china/2026/03/18/is-cheap-energy-the-key-to-china-gaining-ai-supremacy>.
28. Joe Leahy et al., “Robot Nation: China’s Bid to Beat Its Demographic Decline,” *Financial Times*, June 25, 2026, <https://www.ft.com/content/c8731833-10ca-4a12-bfe4-8ebb2584ec68?syn-25a6b1a6=1>; Joshua Yang and Amaya Verde, “China’s Demographic Crisis Means It’s Going to Run Out of Workers,” *Washington Post*, October 11, 2025, <https://www.washingtonpost.com/world/interactive/2025/china-population-decline-worker-crisis/>; and Eleanor Olcott, “China Registers Lowest Number of Births Since Records Began,” *Financial Times*, January 18, 2026, <https://www.ft.com/content/a294436d-9d74-49f5-873e-22ba7a83365b>.
29. Matt Sheehan, “China Is Getting Worried About AI and Jobs,” *Matt Sheehan’s Newsletter*, April 6, 2026, <https://mattsheehan.substack.com/p/china-is-getting-worried-about-ai>; Rachel Cheung, “Balancing AI Pains and Gains,” *The Wire China*, June 7, 2026, <https://www.thewirechina.com/2026/06/07/balancing-ai-pains-and-gains>.
30. Catie Edmondson, “China Wants AI to Flourish, but Not at the Expense of Jobs,” *New York Times*, May 19, 2026, <https://www.nytimes.com/2026/05/19/business/china-ai-unemployment.html>; Jennifer Pak, “A Tech Worker in China Is Laid off and Replaced by AI: Is It Legal?” NPR, May 1, 2026, <https://www.npr.org/2026/05/01/nx-sl-5807131/tech-worker-china-ai>; and Laurie Chen, “China Inc Deploys ‘Quiet’ Layoffs as Beijing Promotes AI Adoption,” Reuters, June 10, 2026, <https://www.reuters.com/business/world-at-work/china-inc-deploys-quiet-layoffs-beijing-promotes-ai-adoption-2026-06-10>.
31. Fergus Ryan et al., *The Party’s AI: How China’s New AI Systems Are Reshaping Human Rights* (Australian Strategic Policy Institute, December 2025), <https://www.aspi.org.au/report/the-partys-ai-how-chinas-new-ai-systems-are-reshaping-human-rights/>; Josh Chin and Liza Lin, *Surveillance State: Inside China’s Quest to Launch a New Era of Social Control* (St. Martin’s Press, 2022).
32. Eleanor Olcott, “China Overhauls World’s Biggest Surveillance Network with Advanced AI,” *Financial Times*, May 27, 2026, <https://www.ft.com/content/f8fa4739-4359-4720-af77-9be1e8370f82>; Julian E. Barnes, “China Aims AI at Predicting Who Could Pose a Political Risk,” *New York Times*, June 1, 2026, <https://www.nytimes.com/2026/06/01/us/politics/china-ai-predicting-dissent.html>; and De Zheng, “New Chinese Surveillance Leaves Foreigners Nowhere to Hide,” DW, May 24, 2026, <https://www.dw.com/en/new-chinese-surveillance-leaves-foreigners-nowhere-to-hide/a-77246713>.
33. Vincent Brussee, “China’s Social Credit System Is Actually Quite Boring,” *Foreign Policy*, September 15, 2021, <https://foreignpolicy.com/2021/09/15/china-social-credit-system-authoritarian>; Chuncheng Liu and Akos Rona-Tas, “Trusting by Numbers: An Analysis of a Chinese Social Credit System Governance Infrastructure,” *Critical Sociology* (2024), summarized by Stanford Center on China’s Economy and Institutions, “Assessing China’s ‘National Model’ Social Credit System,”

- August 2025, <https://scei.fsi.stanford.edu/china-briefs/assessing-chinas-national-model-social-credit-system>.
34. White & Case, “AI Watch: Global Regulatory Tracker – China,” White & Case, September 22, 2025, <https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-china>; Vincent Chow, “Beijing Mandates Internal AI Ethics Reviews to Ensure ‘Controllable’ Tech,” *South China Morning Post*, April 3, 2026, <https://www.scmp.com/tech/policy/article/3348995/beijing-mandates-internal-ai-ethics-reviews-ensure-controllable-tech>; Du Zhihang, “China Issues Guidelines to Standardize AI Agent Development,” *Caixin Global*, May 12, 2026, <https://www.caixinglobal.com/2026-05-12/china-issues-guidelines-to-standardize-ai-agent-development-102443370.html>; and Zilan Qian, “Labelling AI-Generated Content in China,” *Oxford China Policy Lab* (Substack), June 26, 2026, <https://ocpl.substack.com/p/labelling-ai-generated-content-in>.
35. For a fulsome analysis of military AI in the context of U.S.-China competition, see Jacob Stokes et al., *U.S.-China Competition and Military AI: How Washington Can Manage Strategic Risks amid Rivalry with Beijing* (Center for a New American Security, July 25, 2023), <https://www.cnas.org/publications/reports/u-s-china-competition-and-military-ai>; Austin Horng-En Wang et al., *The People’s Liberation Army’s Perspectives on Artificial Intelligence Highlighting Integration as Key to “Intelligentization” Goals* (RAND Corporation, March 2026), <https://www.rand.org/pubs/perspectives/PEA4574-1.html>.
36. The official translation calls this a “green channel,” while other analysts have translated the idea as a “fast-track channel.” “Implementation of the 2025 Plan for National Economic and Social Development and on the 2026 Draft Plan for National Economic and Social Development,” *Xinhua*, March 14, 2026, <https://english.news.cn/20260314/66c5129238d149a08361dbdf53c2431e/c.html>; Jane Tang (@ccjanetang), “NEW: China has 6,200+ AI firms. That’s the latest official count. Its new Five-Year Plan creates a ‘fast-track channel’ pushing commercial technology into military hands—and many companies are moving there willingly. Hangzhou-based Jingan Technology is one of them,” *X* (formerly Twitter), March 26, 2026, <https://x.com/ccjanetang/status/2037160417378640311>.
37. Wang Yi, “Implementing the Global Governance Initiative for a Community with a Shared Future for Humanity,” speech, Lanting Forum, Beijing, October 27, 2025, Ministry of Foreign Affairs of the People’s Republic of China, https://www.fmprc.gov.cn/eng/wjb/wjbz/jh/202510/t20251027_11741537.html; Xi Jinping, “Jointly Forging a Sustainable and Brighter Future,” speech, 32nd APEC Economic Leaders’ Meeting, Gyeongju, November 1, 2025, Ministry of Foreign Affairs of the People’s Republic of China, https://www.fmprc.gov.cn/eng/zy/jj/xjpcx-32apecbdghjxgfw/202511/t20251101_11745402.html.
38. Keith Bradsher, “China’s Leader, With Rare Bluntness, Blames U.S. Containment for Troubles,” *New York Times*, March 7, 2023, <https://www.nytimes.com/2023/03/07/world/asia/china-us-xi-jinping.html>.
39. Lu Chuanying and Cai Yue, “The New Phase of AI Rivalry Between China and the US in the Trump 2.0 Era,” *Pacific Journal*, December 18, 2025, translated by David Cowhig, 高大伟 *David Cowhig’s Translation Blog*, February 28, 2026, <https://gaodawei.wordpress.com/2026/02/28/2025-prc-scholars-on-sino-us-trump-2-0-ai-rivalry>.
40. “Four Great Inventions of China,” Embassy of the People’s Republic of China in Antigua and Barbuda, November 12, 2013, https://ag.china-embassy.gov.cn/eng/zgyabjw/202302/t20230210_11023529.htm.
41. Chen Na, “Chinese Museum Commemorating ‘Two Bombs, One Satellite’ Reopens,” *Chinese Academy of Sciences*, October 18, 2024, https://english.cas.cn/newsroom/cas_media/202410/t20241018_691957.shtml; Yang Sheng, “People Hail ‘Two Bombs, One Satellite’ Founders for Inspiring Generations of Chinese Scientists,” *Global Times*, October 16, 2019, <https://www.globaltimes.cn/content/1167113.shtml>.
42. William Hannas et al., *China’s Advanced AI Research: Monitoring China’s Paths to “General” Artificial Intelligence* (Center for Security and Emerging Technology, July 2022), <https://cset.georgetown.edu/publication/chinas-advanced-ai-research>; Zilan Qian, “How China Hopes to Build AGI Through Self-Improvement,” *ChinaTalk*, March 30, 2026, <https://www.china-talk.media/p/how-china-hopes-to-build-agi-through>.
43. Kyle Chan, “Does China Care about AGI?,” *High Capacity* (Substack), February 15, 2026, <https://www.high-capacity.com/p/does-china-care-about-agi>.
44. John Liu, “China’s AI Industry Looks Unstoppable in the Race to Best US Rivals, but Is It?” *CNN*, February 11, 2026, <https://www.cnn.com/2026/02/10/tech/china-us-ai-race-challenges-intl-hnk-dst>.
45. Schaefer and Nunlist, “The AI Plus Initiative: China’s Blueprint for AI Diffusion.”
46. Jonathan Cheng, “China Is Worried AI Threatens Party Rule—and Is Trying to Tame It,” *Wall Street Journal*, December 24, 2025, <https://www.wsj.com/tech/ai/china-is-worried-ai-threatens-party-ruleand-is-trying-to-tame-it-bfdcd2d>.
47. Brian Tse, “China Is Taking AI Safety Seriously: So Must the U.S.,” *Time*, August 13, 2025, <https://time.com/7308857/china-isnt-ignoring-ai-regulation-the-u-s-shouldnt-either>; Chow, “Beijing Mandates Internal AI Ethics Reviews to Ensure ‘Controllable’ Tech.”
48. “省部级主要领导干部专题研讨班侧记 [Notes on the Seminar for Provincial and Ministerial-Level Leading Cadres],” *Xinhua News Agency*, January 24, 2026, <https://www.news.cn/politics/leaders/20260124/45b786058d7d4f7c9ef9578380c2bd9d/c.html>.
49. Ben Jiang, “China’s Zhipu AI Ramps Up Push Towards Human-Level AI, Reiterates Open Source Pledge,” *South China Morning Post*, December 26, 2025, <https://www.scmp.com/tech/big-tech/article/3337822/chinas-zhipu-ai-ramps-push-towards-human-level-ai-reiterates-open-source-pledge>.

50. Evelyn Cheng and Annie Bao, “China Poaches More AI Talent from the U.S. as It Eyes the Next ‘Super-App,’” *CNBC*, June 5, 2026, <https://www.cnbc.com/2026/06/05/china-may-move-toward-us-path-on-ai-as-firms-poach-employees.html>.
51. Reuters, “Chinese AI Startup Z.ai Closes Frontier Gap with US Rivals, Plans Shanghai Listing,” *CNBC TV18*, June 25, 2026, <https://www.cnbctv18.com/technology/chinese-startup-zai-closes-frontier-gap-with-anthropic-plans-shanghai-listing-19932417.htm>; Laurie Chen, “After Anthropic Shutdown, China’s Z.ai Closes Frontier Gap As It Plans Dual Listing,” *Reuters*, June 25, 2026, <https://www.reuters.com/world/asia-pacific/after-anthropic-shutdown-chinas-zai-closes-frontier-gap-it-plans-dual-listing-2026-06-25>.
52. Xiaofei Xu and Wency Chen, “DeepSeek Hiring Spree: Chinese AI Firm Seeks Newcomers as It Pursues AGI,” *South China Morning Post*, June 26, 2026, <https://www.scmp.com/tech/big-tech/article/3358394/deepseek-hiring-spree-chinese-ai-firm-seeks-newcomers-it-pursues-agi>.
53. “About Us,” Beijing Institute for General Artificial Intelligence, accessed May 15, 2026, <https://eng.bigai.ai/about/>; Huey-Meei Chang and William Hannas, *Spotlight on Beijing Institute for General Artificial Intelligence: China’s State-Backed Program for General Purpose AI* (Center for Security and Emerging Technology, May 2023), <https://cset.georgetown.edu/publication/spotlight-on-beijing-institute-for-general-artificial-intelligence/>.
54. Xinhua, “Chinese Researchers Score Breakthrough in General AI Logical Reasoning,” *China Daily Asia*, January 27, 2026, <https://www.chinadailyasia.com/hk/article/627883>.
55. Wagner, “AI Safety in the 15th Five-Year Plan—Two Sessions Special Edition.”
56. *China’s Drive toward Self-Reliance in Artificial Intelligence: From Chips to Large Language Models* (Mercator Institute for China Studies, July 22, 2025), <https://merics.org/en/report/chinas-drive-toward-self-reliance-artificial-intelligence-chips-large-language-models>.
57. Sunny Cheung, “China’s New Five-Year Plan Prioritizes Robotics. The World Should Pay Attention,” *The Diplomat*, March 14, 2026, <https://thediplomat.com/2026/03/chinas-new-five-year-plan-prioritizes-robotics-the-world-should-pay-attention>.
58. “China’s First National Standard System for Humanoid Robotics Poised to Spur Industry Development,” *Xinhua*, March 3, 2026, <https://english.news.cn/20260303/0e51ac8f-66c542c5bacf2af3f80b3a40/c.html>. For analysis of China’s approach to robotics and embodied AI, see Pavlo Zvenyhorodskyi 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?lang=en>; William Hannas et al., *China’s Embodied AI: A Path to AGI* (Center for Security and Emerging Technology, December 2025), <https://cset.georgetown.edu/publication/chinas-embodied-ai-a-path-to-agi>; and Meaghan Tobin and Xinyun Wu, “Does China Have a Robot Bubble?,” *New York Times*, December 17, 2025, <https://www.nytimes.com/2025/12/17/business/china-robot-unitree-ubtech.html>. For an analysis of different approaches to achieving AGI, see William Marcellino et al., *Charting Multiple Courses to Artificial General Intelligence* (RAND, April 23, 2025), <https://www.rand.org/pubs/perspectives/PEA3691-1.html>. For an analysis of Chinese labs’ thinking on world models, see “AI and DeepSeek in Kazakhstan: Berkeley Investigates Qwen ‘Lying’; Elite Scientists Forecast 2026,” *Lingua Sinica* (Substack), January 16, 2026, <https://linguasinica.substack.com/p/01ai-and-deepseek-in-kazakhstan-berkeley>. And for a primer on world model types, see Chris Paxton, “Will World Models Allow Robots to Think?” *It Can Think!* (Substack), March 26, 2026, <https://itcantthink.substack.com/p/will-world-models-allow-robots-to>.
59. Meghan Bobrowsky, “He’s Been Right About AI for 40 Years. Now He Thinks Everyone Is Wrong,” *Wall Street Journal*, November 14, 2025, <https://www.wsj.com/tech/ai/yann-lecun-ai-meta-0058b13c>; Cade Metz, “An A.I. Pioneer Warns the Tech ‘Herd’ Is Marching Into a Dead End,” *New York Times*, January 26, 2026, <https://www.nytimes.com/2026/01/26/technology/an-ai-pioneer-warns-the-tech-herd-is-marching-into-a-dead-end.html>.
60. Fei-Fei Li, “Spatial Intelligence Is AI’s Next Frontier,” *Time*, December 11, 2025, <https://time.com/7339693/fei-fei-li-ai>.
61. U.S.-China Economic and Security Review Commission, “Comprehensive List of the Commission’s 2024 Recommendations,” November 2024, https://www.uscc.gov/sites/default/files/2024-11/2024_Comprehensive_List_of_Recommendations.pdf; U.S. Department of Energy, “Secretary Wright Leads AI Collaboration Event from Oak Ridge National Lab,” *Energy.gov*, February 28, 2025, <https://www.energy.gov/articles/secretary-wright-leads-ai-collaboration-event-oak-ridge-national-lab>.
62. “Made in China 2025 and China’s Industrial Policies: Issues for Congress,” IF10964, Congressional Research Service, updated December 12, 2024, <https://www.congress.gov/crs-product/IF10964>; Camille Boullenois et al., *Was Made in China 2025 Successful?* (Rhodium Group, May 5, 2025), <https://rhg.com/research/was-made-in-china-2025-successful>.
63. Joseph Webster, “Beijing Military City May Be Top Cover for a Covert AGI Program,” *China-Russia Report* (Substack), August 14, 2025, <https://chinarussiareport.substack.com/p/beijing-military-city-may-be-top>; Joseph Webster, “Sprinters, Marathoners & Skeptics on the Future of AI & Power,” *War on the Rocks*, August 1, 2025, <https://warontherocks.com/2025/08/sprinters-marathoners-skeptics-on-the-future-of-ai-power>; and Demetri Sevastopulo et al., “China Builds Huge Wartime Military Command Centre in Beijing,” *Financial Times*, January 30, 2025, <https://www.ft.com/content/f3763e51-8607-42b9-9ef9-5789d5bf353d>.
64. Ronan Farrow and Andrew Marantz, “Sam Altman May Control Our Future—Can He Be Trusted?” *New Yorker*, April 13, 2026, <https://www.newyorker.com/magazine/2026/04/13/sam-altman-may-control-our-future-can-he-be-trusted>.

65. "OpenAI Charter," OpenAI, April 2018, <https://openai.com/charter>; Sam Altman, "Reflections," *Sam Altman* (blog), January 6, 2025, <https://blog.samaltman.com/reflections>.
66. Zuckerberg, "Personal Superintelligence," Meta, July 30, 2025, <https://www.meta.com/superintelligence>.
67. Isabel Juniewicz, "Hyperscaler Capex Has Quadrupled since GPT-4's Release," Epoch AI, February 26, 2026, <https://epoch.ai/data-insights/hyperscaler-capex-trend>.
68. "Announcing the Stargate Project," OpenAI, January 21, 2025, <https://openai.com/index/announcing-the-stargate-project>.
69. *Winning the Race: America's AI Action Plan* (The White House, July 23, 2025), 1, <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.
70. Ina Fried and Madison Mills, "AI Spending Is Eclipsing Salaries at Some Companies, and Execs Are Starting to Notice," Axios AI+, April 27, 2026, <https://www.axios.com/newsletters/axios-ai-plus-907e3f10-3ffb-11f1-bf3a-23dcada653e3.html>.
71. "What If AI Made the World's Economic Growth Explode?" *Economist*, July 24, 2025, <https://www.economist.com/briefing/2025/07/24/what-if-ai-made-the-worlds-economic-growth-explode>.
72. Bryan Walsh, "We're Running out of Good Ideas. AI Might Be How We Find New Ones," Vox, December 13, 2025, <https://www.vox.com/future-perfect/471918/ai-science-growth-deep-mind-alphafold-chatgpt-google>.
73. Ina Fried, "Models That Improve on Their Own Are AI's Next Big Thing," Axios, January 27, 2026, <https://www.axios.com/2026/01/27/models-improve-ai>; Helen Toner et al., *When AI Builds AI: Findings From a Workshop on Automation of AI R&D* (Center for Security and Emerging Technology, January 2026), <https://cset.georgetown.edu/publication/when-ai-builds-ai>.
74. Frank Flight, "The 2026 Global Intelligence Crisis," Citadel Securities, February 24, 2026, <https://www.citadelsecurities.com/news-and-insights/2026-global-intelligence-crisis>.
75. Erik R. Quam, "PRC Conceptions of Comprehensive National Power: Part 1," *China Brief* 25, no. 16 (2025), <https://jamestown.org/prc-conceptions-of-comprehensive-national-power-part-1>.
76. Daniel Thilo Schroeder et al., "How Malicious AI Swarms Can Threaten Democracy," *Science* 391, no. 6783 (January 22, 2026): 354–357, <https://doi.org/10.1126/science.adz1697>; David Gilbert, "AI-Powered Disinformation Swarms Are Coming for Democracy," *Wired*, January 22, 2026, <https://www.wired.com/story/ai-powered-disinformation-swarms-are-coming-for-democracy>.
77. "Chinese Artificial Intelligence Distorts Perceptions," in *Annual Report 2026* (Estonian Foreign Intelligence Service, 2026), <https://raport.valisluureamet.ee/2026/en/6-asia/6-3-chinese-artificial-intelligence-distorts-perceptions>.
78. Stella Chen, "The CMP Dictionary: Discourse Power," China Media Project, May 30, 2022, https://chinamediaproject.org/the_ccp_dictionary/discourse-power.
79. Cristina Criddle, "Microsoft Warns That China Is Winning AI Race Outside the West," *Financial Times*, January 13, 2026, <https://www.ft.com/content/f7a5b184-1fef-4f02-b957-4c2b07adf91f>; Sarah Cook, "China's DeepSeek: Implications of Its Growing Reach in Autocracies and Democracies," *UnderReported China* (Substack), January 30, 2026, <https://underreportedchina.substack.com/p/chinas-deepseek-implications-for-democracies-and-autocracies>.
80. Steven Lee Myers, "DeepSeek's Answers Include Chinese Propaganda, Researchers Say," *New York Times*, January 31, 2025, <https://www.nytimes.com/2025/01/31/technology/deepseek-chinese-propaganda.html>; Jinghan Zeng, "ChatGPT as a Security Threat: US-China Security Dilemma in the Generative AI Race," *The British Journal of Politics and International Relations* (2025), <https://journals.sagepub.com/doi/10.1177/13691481251389415>.
81. Hannah Waight et al., "State Media Control Influences Large Language Models," *Nature* (2026), <https://state-media-influence-llm.github.io>; Lingling Wei, "The Hidden Chinese Influence in AI," *Wall Street Journal*, May 19, 2026, <https://www.wsj.com/world/china/the-hidden-chinese-influence-in-ai-c2837047>.
82. Ryan et al., *The Party's AI: How China's New AI Systems Are Reshaping Human Rights*; and David Bandurski, "AI for Human Propaganda," *Lingua Sinica* (Substack), April 2, 2026, <https://linguasinica.substack.com/p/ai-for-human-propaganda>.
83. Xinhua, "President Xi Urges New Media Outlet to 'Tell China Stories Well,'" Xinhua, December 31, 2016, http://www.xinhuanet.com/english/2016-12/31/c_135946550.htm; "The CMP Dictionary: Telling China's Story Well," China Media Project, April 16, 2021, https://chinamediaproject.org/the_ccp_dictionary/telling-chinas-story-well.
84. Alexander Boyd, "The Historical Nihil-List: Cyberspace Administration Targets Top Ten Deviations from Approved History," *China Digital Times*, August 16, 2021, <https://chinadigitaltimes.net/2021/08/the-historical-nihil-list-cyberspace-administration-targets-top-ten-deviations-from-approved-history>.
85. Theo Burman, "How China Is Training AI to Censor Its Secrets," *Newsweek*, March 28, 2025, <https://www.newsweek.com/china-ai-training-censorship-llm-2052117>; OpenAI, "Disrupting Malicious Uses of Our Models: An Update, February 2026," OpenAI, February 2026, <https://cdn.openai.com/pdf/df438d70-e3fe-4a6c-a403-ff632def8f79/disrupting-malicious-uses-of-ai.pdf>; and Sean Lyngaas, "A Chinese Official's Use of ChatGPT Accidentally Revealed a Global Intimidation Operation," CNN, February 25, 2026, <https://www.cnn.com/2026/02/25/politics/chatgpt-china-intimidation-operation>.

86. Joseph S. Nye, Jr., *Soft Power: The Means to Success in World Politics* (Public Affairs Books, 2005).
87. Youlun Nie, “New PRC Cybercrime Law Heralds Digital Iron Curtain,” *China Brief* 26, no. 5 (March 6, 2026), <https://jamestown.org/new-prc-cybercrime-law-heralds-digital-iron-curtain>.
88. “China Is Using Facebook to Build a Huge Audience Around the World,” *Economist*, April 20, 2019, <https://www.economist.com/graphic-detail/2019/04/20/china-is-using-facebook-to-build-a-huge-audience-around-the-world>.
89. Five Eyes, “Five Eyes Cyber Security Agencies Statement,” press release, June 22, 2026, <https://www.ncsc.gov.uk/sites/default/files/2026-06/Five-Eyes-cyber-security-agencies-statement-ai-shift.pdf>; Kevin Roose, “Anthropic Claims Its New AI Model, Mythos, Is a Cybersecurity ‘Reckoning,’” *New York Times*, April 7, 2026, <https://www.nytimes.com/2026/04/07/technology/anthropic-claims-its-new-ai-model-mythos-is-a-cybersecurity-reckoning.html>; and Anthropic, “Assessing Claude Mythos Preview’s Cybersecurity Capabilities,” Anthropic, April 7, 2026, <https://red.anthropic.com/2026/mythos-preview>.
90. Cristina Criddle et al., “Pentagon Moves to Build AI Tools for China Cyber Operations,” *Financial Times*, February 27, 2026, <https://www.ft.com/content/a56d70b5-669c-4bcc-8541-a4961fc99802>; Maxwell Zeff, “Trump Administration Allows Anthropic to Release Mythos to Select US Organizations,” *Wired*, June 26, 2026, <https://www.wired.com/story/anthropic-restores-access-to-mythos/>; and Cheyenne Haslett, “Trump Administration Steps in to Limit OpenAI’s Latest Model Launch,” *Politico*, June 26, 2026, <https://www.politico.com/news/2026/06/25/openai-gpt-model-government-approval-00977551>.
91. “Baidu Robotaxis Stranded on Roads in Central China, Wuhan Due to ‘System Failure,’ Police Say,” CNN, April 2, 2026, <https://www.cnn.com/2026/04/02/china/baidu-robotaxi-wuhan-system-failure-intl-hnk>.
92. “A History of Biological Weapons,” PBS American Experience, accessed May 18, 2026, <https://www.pbs.org/wgbh/american-experience/features/weapon-timeline>.
93. Adrian Woolfson, *On the Future of Species: Authoring Life by Means of Artificial Biological Intelligence* (MIT Press, April 28, 2026), <https://mitpress.mit.edu/9780262054898/on-the-future-of-species>; Georgia Adamson and Gregory C. Allen, “Opportunities to Strengthen U.S. Biosecurity from AI-Enabled Bioterrorism: What Policymakers Should Know,” Center for Strategic and International Studies, August 6, 2025, <https://www.csis.org/analysis/opportunities-strengthen-us-biosecurity-ai-enabled-bioterrorism-what-policymakers-should>.
94. Jacob Stokes et al., *Averting AI Armageddon: U.S.-China-Russia Rivalry at the Nexus of Nuclear Weapons and Artificial Intelligence* (Center for a New American Security, February 13, 2025), <https://www.cnas.org/publications/reports/averting-ai-armageddon>; Simon Goldstein and Peter N. Salib, “Nuclear Deterrence in the Age of AGI,” *Lawfare*, March 26, 2025, <https://www.lawfaremedia.org/article/nuclear-deterrence-in-the-age-of-agi>.
95. Richard Hundley, *Past Revolutions, Future Transformations: What Can the History of Revolutions in Military Affairs Tell Us About Transforming the U.S. Military?* (RAND Corporation, 1999), https://www.rand.org/pubs/monograph_reports/MR1029.html.
96. Yoshua Bengio et al., *International AI Safety Report 2026* (International Scientific Report on the Safety of Advanced AI, February 3, 2026), 72, <https://internationalaisafetyreport.org/sites/default/files/2026-02/international-ai-safety-report-2026.pdf>.
97. Michael J. D. Vermeer, *Evaluating Select Global Technical Options for Countering a Rogue AI* (RAND Corporation, November 2025), <https://www.rand.org/pubs/perspectives/PEA4361-1.html>.
98. Monte MacDiarmid et al., “Natural Emergent Misalignment from Reward Hacking in Production RL,” Anthropic, November 2025, <https://assets.anthropic.com/m/74342f2c96095771/original/Natural-emergent-misalignment-from-reward-hacking-paper.pdf>; Matthias Bastian, “Strict Anti-Hacking Prompts Make AI Models More Likely to Sabotage and Lie, Anthropic Finds,” *The Decoder*, November 23, 2025, <https://the-decoder.com/strict-anti-hacking-prompts-make-ai-models-more-likely-to-sabotage-and-lie-anthropic-finds>.
99. Will Knight, “AI Models Lie, Cheat, and Steal to Protect Other Models from Being Deleted,” *Wired*, April 1, 2026, <https://www.wired.com/story/ai-models-lie-cheat-steal-protect-other-models-research>.
100. See, among others, Arianna Huffington, “The Human-AI Alignment Problem,” *Time*, November 23, 2025, <https://time.com/7335983/human-ai-alignment-problem>.
101. “Declassified Assessment on COVID-19 Origins,” Office of the Director of National Intelligence, October 29, 2021, <https://www.dni.gov/files/ODNI/documents/assessments/Declassified-Assessment-on-COVID-19-Origins.pdf>.
102. “U.S.-China Climate Cooperation Timeline,” Asia Society Policy Institute, accessed May 18, 2026, <https://asiasociety.org/policy-institute/us-china-climate-cooperation-timeline>; Andrew S. Erickson and Gabriel Collins, “Competition with China Can Save the Planet,” *Foreign Affairs*, April 13, 2021, <https://www.foreignaffairs.com/articles/united-states/2021-04-13/competition-china-can-save-planet>.
103. John Cassidy, “The Dangerous Paradox of A.I. Abundance,” *New Yorker*, January 12, 2026, <https://www.newyorker.com/news/the-financial-page/the-dangerous-paradox-of-ai-abundance>; Annie Lowrey, “The Worst-Case Future for White-Collar Workers,” *Atlantic*, February 18, 2026, <https://www.theatlantic.com/ideas/2026/02/ai-white-collar-jobs/686031>; and Gina Raimondo, “America Cannot Withstand the Economic Shock That’s Coming,” *New York Times*, March 6, 2026, <https://www.nytimes.com/2026/03/06/opinion/ai-labor-unemployment.html>.

104. Marietje Schaake, “The ‘Botlash’ Movement Is Gaining Momentum,” *Financial Times*, February 24, 2026, <https://www.ft.com/content/ecead6b9-eb42-4a85-bd33-073c659e84bf>; Rebecca Lissner, “A Populist Backlash over AI Is Brewing in America,” *Time*, February 6, 2026, <https://time.com/7371825/trump-data-center-ai-backlash-ai-america-china>; Matthew Daly, “Sanders, Ocasio-Cortez Push Bill to Impose AI Data Center Moratorium,” Associated Press, March 25, 2026, <https://apnews.com/article/data-centers-ai-electricity-sanders-aoc-65651bd28c3d911d18eeb46cd54f4c75>; and Jason Lange and Courtney Rozen, “Half of Americans Fear AI Could Put Someone in Their Household Out of Work, Reuters/Ipsos Poll Finds,” Reuters, June 10, 2026, <https://www.reuters.com/business/world-at-work/half-americans-fear-ai-could-put-someone-their-household-out-work-reutersipsos-2026-06-10>.
105. Eric Platt, “BlackRock’s Larry Fink Warns AI May Intensify Wealth Inequality,” *Financial Times*, March 23, 2026, <https://www.ft.com/content/58602814-8138-4a14-a32e-d82dec2d472c>.
106. Daron Acemoglu et al., *AI, Human Cognition and Knowledge Collapse*, working paper (National Bureau of Economic Research, February 2026), <https://www.nber.org/papers/w34910>.
107. Sheehan, “China Is Getting Worried About AI & Jobs.”
108. Barclay Bram, “The 19 Percent Revisited: How Youth Unemployment Has Changed Chinese Society,” Asia Society Policy Institute, September 3, 2025, <https://asiasociety.org/policy-institute/19-percent-revisited-how-youth-unemployment-has-changed-chinese-society>.
109. Lingling Wei, “China’s New Strategy for Trump: Punch Hard, Concede Little,” *Wall Street Journal*, October 23, 2025, <https://www.wsj.com/world/china/china-trump-strategy-06841606>; “Tencent Vows Fresh Gaming Curbs After ‘Spiritual Opium’ Attack Zaps \$60 Billion,” Reuters, August 3, 2021, <https://www.reuters.com/technology/tencent-falls-after-china-media-calls-online-gaming-spiritual-opium-2021-08-03>.
110. Daniel Ten Kate, “China’s AI Nightmare Is an Out-of-Control Welfare State,” Bloomberg, March 6, 2026, <https://www.bloomberg.com/news/features/2026-03-06/will-ai-disruption-force-xi-jinping-to-turn-china-into-a-welfare-state>.
111. Jed Kolko, “Research on AI and the Labor Market Is Still in the First Inning,” Peterson Institute for International Economics, March 10, 2026, <https://www.piie.com/blogs/realtime-economics/2026/research-ai-and-labor-market-still-first-inning>; Alex Martin Richmond, *The AI Jobs Transition Framework: Mapping AI’s Near-Term Impact on Jobs* (OpenAI, April 2026), <https://cdn.openai.com/pdf/the-ai-jobs-transition-framework-report.pdf>.
112. *The Next Great Divergence: Why AI May Deepen Inequality between Countries* (United Nations Development Programme, December 2025), <https://www.undp.org/asia-pacific/next-great-divergence>.
113. Cristina Criddle et al., “Rich Countries’ Greater Use of AI Risks Deepening Inequality, Anthropic warns,” *Financial Times*, January 15, 2026, <https://www.ft.com/content/3ad44e30-c738-4356-91fb-8bb2368685c4>; Ruth Appel et al., *The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption* (Anthropic, September 15, 2025), <https://www-cdn.anthropic.com/7b76335c444876a93fa22a63aabb4aeb820aff25.pdf>.
114. Muath Alduhishy, “Sovereign AI: What It Is, and 6 Strategic Pillars for Achieving It,” World Economic Forum, April 25, 2024, <https://www.weforum.org/stories/2024/04/sovereign-ai-what-is-ways-states-building>; Juan N. Pava et al., “AI Sovereignty’s Definitional Dilemma,” Stanford HAI, February 17, 2026, <https://hai.stanford.edu/news/ai-sovereignty-s-definitional-dilemma>; and Pablo Chavez et al., *Sovereign AI Index: Tracking the Global Push for AI Self-Reliance* (CNAS, April 2026), <https://interactives.cnas.org/reports/sovereign-ai-index>.
115. “Read Mark Carney’s Full Speech on Middle Powers Navigating a Rapidly Changing World,” CBC News, January 20, 2026, <https://www.cbc.ca/news/politics/mark-carney-speech-davos-rules-based-order-9.7053350>.
116. Jacob Poushter et al., *How People Around the World View AI* (Pew Research Center, October 15, 2025), <https://www.pewresearch.org/global/2025/10/15/how-people-around-the-world-view-ai>.
117. Mark F. Cancian, *Avoiding Coping with Surprise in Great Power Conflicts* (CSIS, 2018), 154, <https://www.jstor.org/stable/resrep22428>. Analysts have argued the next surprise could come in robotics and embodied AI; see Scott Singer and Pavlo Zvenyhorodskiy, “DeepSeek Was a Warning Shot. China Is Building Its Next Surprise,” *Washington Post*, January 30, 2026, <https://www.washingtonpost.com/opinions/2026/01/30/china-ai-robots-autonomous-drones>; Maria Langan Reikhs and Herb Lin, *Anticipating and Navigating Strategic Technological Surprise* (Hoover Institution, May 26, 2026), https://hoover-s3-website.s3.us-west-2.amazonaws.com/s3fs-public/research/docs/Strategic%20Surprise_5-26-26.pdf.
118. Joe Khawam, “The Case for Imposing Costs on China’s AI Distillation Campaigns,” Just Security, March 30, 2026, <https://www.justsecurity.org/134124/costs-china-ai-distillation>; Shirin Ghaffary and Maggie Eastland, “OpenAI, Anthropic, Google Unite to Combat Model Copying in China,” Bloomberg, April 6, 2026, <https://www.bloomberg.com/news/articles/2026-04-06/openai-anthropic-google-unite-to-combat-model-copying-in-china>.
119. Executive Office of the President, Office of Science and Technology Policy, “Memorandum for the Heads of Executive Departments and Agencies,” press release, April 23, 2026, <https://whitehouse.gov/wp-content/uploads/2026/04/NSTM-4.pdf>; Daniel Remler and Ben Hayum, *Adversarial Distillation: China’s Campaign to Extract American AI Capabilities* (CNAS, June 2, 2026), <https://www.cnas.org/publications/reports/adversarial-distillation>.
120. Olivia Gaziz, “CIA Director William Burns: ‘I Wouldn’t Underestimate’ Xi’s Ambitions for Taiwan,” CBS News, February 3, 2023, <https://www.cbsnews.com/news/cia-director-william-burns-i-wouldnt-underestimate-xis-ambitions-for-taiwan>.

121. Richard McGregor and Jude Blanchette, “After Xi: Future Scenarios for Leadership Succession in Post-Xi Jinping Era,” Lowy Institute and the Center for Strategic and International Studies, April 21, 2021, <https://www.loyyinstitute.org/publications/after-xi-future-scenarios-leadership-succession-post-xi-jinping-era>.
122. It is theoretically possible that some third country or entity achieves AGI before the United States or China, but it is highly unlikely given the scale needed to train models powerful enough to be AGI and the fact that third countries lack the necessary inputs in sufficient amounts.
123. David E. Sanger, “Obama Order Sped Up Wave of Cyberattacks Against Iran,” *New York Times*, June 1, 2012, <https://www.nytimes.com/2012/06/01/world/middleeast/obama-ordered-wave-of-cyberattacks-against-iran.html>.
124. Apostol Vassilev et al., *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations* (National Institute of Standards and Technology, 2023), <https://csrc.nist.gov/pubs/ai/100/2/e2023/final>; *Understanding Adversarial Attacks Against Machine Learning and AI* (National Cyber Security Centre, April 29, 2026), <https://www.ncsc.gov.uk/paper/understanding-adversarial-attacks-against-machine-learning-and-ai>.
125. U.S. and Chinese forces fought each other during the Korean and Vietnam Wars, but the two sides have never fought each other in a war that included direct attacks on each other’s territories. See “Battle of the Chosin Reservoir,” *Encyclopedia Britannica*, <https://www.britannica.com/event/Battle-of-the-Chosin-Reservoir>; Reuters, “China Admits Combat in Vietnam War,” *Washington Post*, May 17, 1989, <https://www.washingtonpost.com/archive/politics/1989/05/17/china-admits-combat-in-vietnam-war/6b9cb8a4-4d18-48bf-80d2-bea80f64057c>.
126. Paul Scharre and Megan Lamberth, *Artificial Intelligence and Arms Control* (CNAS, October 12, 2022), <https://www.cnas.org/publications/reports/artificial-intelligence-and-arms-control>.
127. The same concept is also sometimes called a “breakout estimate.” See, for example, “The Historic Deal That Will Prevent Iran from Acquiring a Nuclear Weapon,” The White House, <https://obamawhitehouse.archives.gov/node/328996>. And for further analysis of breakout time calculations, see Kelsey Davenport and Julia Masterson, “The Limits of Breakout Estimates in Assessing Iran’s Nuclear Program,” Issue Brief, Arms Control Association, August 4, 2020, <https://www.armscontrol.org/issue-briefs/2020-08/limits-breakout-estimates-assessing-irans-nuclear-program>.
128. Dan Hendrycks et al., *Superintelligence Strategy*, <https://www.nationalsecurity.ai/chapter/deterrence-with-mutual-assured-ai-malfunction-maim#how-to-maintain-a-maim-regime>.
129. Kelvin Chan, “Iranian Strikes on Amazon Data Centers Highlight Industry’s Vulnerability to Physical Disasters,” Reuters, March 3, 2026, <https://apnews.com/article/amazon-aws-data-center-uae-iran-bahrain-71066b0a822c4cfd88b61e-3fe79af917>.
130. U.S. Department of Justice, Office of Public Affairs, “Former Google Engineer Found Guilty of Economic Espionage and Theft of Confidential AI Technology,” press release, January 30, 2026, <https://www.justice.gov/opa/pr/former-google-engineer-found-guilty-economic-espionage-and-theft-confidential-ai-technology>; Shirin Ghaffary and Maggie Eastland, “Open AI Accuses China’s DeepSeek of Stealing AI Technology,” *Los Angeles Times*, February 13, 2026, <https://www.latimes.com/business/story/2026-02-13/openai-accuses-chinas-deepseek-of-stealing-ai-technology>.
131. For overviews of how recursive self-improvement in AI systems could happen, see Jack Clark, “Import AI 455: AI Systems Are About to Start Building Themselves,” *Import AI* (Substack), May 4, 2026, <https://importai.substack.com/p/import-ai-455-automating-ai-research>; “How Artificial Intelligence Got Better at Building Itself,” *Economist*, June 7, 2026, <https://www.economist.com/science-and-technology/2026/06/07/how-artificial-intelligence-got-better-at-building-itself>.
132. Christopher Klein, “The Spies Who Launched America’s Industrial Revolution,” *History*, January 10, 2019, <https://www.history.com/articles/industrial-revolution-spies-europe>.
133. Will Knight, “The Humanoid Robot of the Future Is a 6-Foot-Tall Beefcake with a Chinese Body and an American Brain,” *Wired*, June 3, 2026, <https://www.wired.com/story/nvidia-uni-tree-humanoid-robot-h2-plus>.
134. Stokes et al., *Averting AI Armageddon*; and Michael C. Horowitz, “Artificial Intelligence and the Future of Strategic Stability,” *Texas National Security Review* 9, no. 2 (2026), <https://tnsr.org/roundtable/artificial-intelligence-and-the-future-of-strategic-stability>.
135. In her study on technological strategic surprise, Audrey Kurth Cronin argues that technological revolutions affecting state power are either open or closed. AGI would arguably have attributes of both types. See Audrey Kurth Cronin, *Open Innovation* (U.S. Army War College Press, August 14, 2020), <https://press.armywarcollege.edu/cgi/viewcontent.cgi?article=2675&context=parameters>.
136. Khawam, “The Case for Imposing Costs on China’s AI Distillation Campaigns”; Ghaffary and Eastland, “OpenAI, Anthropic, Google Unite to Combat Model Copying in China”; and Executive Office of the President, Office of Science and Technology Policy, “Memorandum for the Heads of Executive Departments and Agencies.” To understand how various entities use a slew of channels to supply the PLA with advanced chips, see John Costello, *How the PLA Continued Procuring Nvidia AI Chips After U.S. Export Controls* (WireScreen, June 1, 2026), <https://www.wirescreen.ai/research/gpu-procurement-report>; Savannah Billman, “How the PLA Has Tried to Get Around Chip Export Control,” *The Wire China*, June 4, 2026, <https://www.thewirechina.com/2026/06/04/how-the-pla-has-tried-to-get-around-chip-export-controls>.
137. Rebecca Tan and Lyric Li, “Beijing Tightens Its Grip on AI Firms That Try to Shed Their Chinese Ties,” *Washington Post*, April 21, 2026, <https://www.washingtonpost.com/>

[world/2026/04/21/china-ai-competition-manus-meta](https://www.wsj.com/world/china/china-bans-metas-acquisition-of-manus-on-national-security-grounds-71e10c3f); Raffaele Huang, “China Bans Meta’s Acquisition of Manus on National Security Grounds,” *Wall Street Journal*, April 27, 2026, <https://www.wsj.com/world/china/china-bans-metas-acquisition-of-manus-on-national-security-grounds-71e10c3f>; and Bloomberg News, “China Expands Travel Curbs to Top AI Talent at Private Firms,” Bloomberg, May 26, 2026, <https://www.bloomberg.com/news/articles/2026-05-26/china-expands-travel-curbs-to-top-ai-talent-at-private-firms>.

138. Maximilian Negele et al., *Europe and the Geopolitics of AGI: The Need for a Preparedness Plan* (RAND Europe, December 16, 2025), https://www.rand.org/pubs/research_reports/RRA4636-1.html.
139. “The National Counterterrorism Center,” Office of the Director of National Intelligence, accessed July 17, 2026, <https://www.dni.gov/index.php/nctc-who-we-are>.
140. Amanda Macias, “FBI Chief Slams Chinese Cyberattacks on U.S., Calls It ‘One of the Largest Transfers of Wealth in Human History,’” CNBC, July 7, 2026, <https://www.cnn.com/2020/07/07/fbi-chief-slams-chinese-cyberattacks-against-us-hudson-institute.html>.
141. Some have called for “AI with a Western soul.” The author instead uses “democratic” to describe systems that are based on and help support democratic values. Democratic regimes now extend far beyond the geographic West, so the latter term can be misleading in the contemporary context. See a16z (@a16z), “@UnderSecPD Sarah B. Rogers: AI with a ‘Western soul’ is the greatest soft power tool America can possess,” X (Twitter), March 4, 2026, <https://x.com/a16z/status/2029255086757843381>.

CNAS Editorial

DIRECTOR OF STUDIES

Katherine L. Kuzminski

PUBLICATIONS & EDITORIAL DIRECTOR

Maura McCarthy

ASSOCIATE EDITOR

Caroline Steel

ASSOCIATE EDITOR

Colleen Kristy

CREATIVE DIRECTOR

Melody Cook

Cover Art & Production Notes

COVER ILLUSTRATION

Melody Cook

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://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

