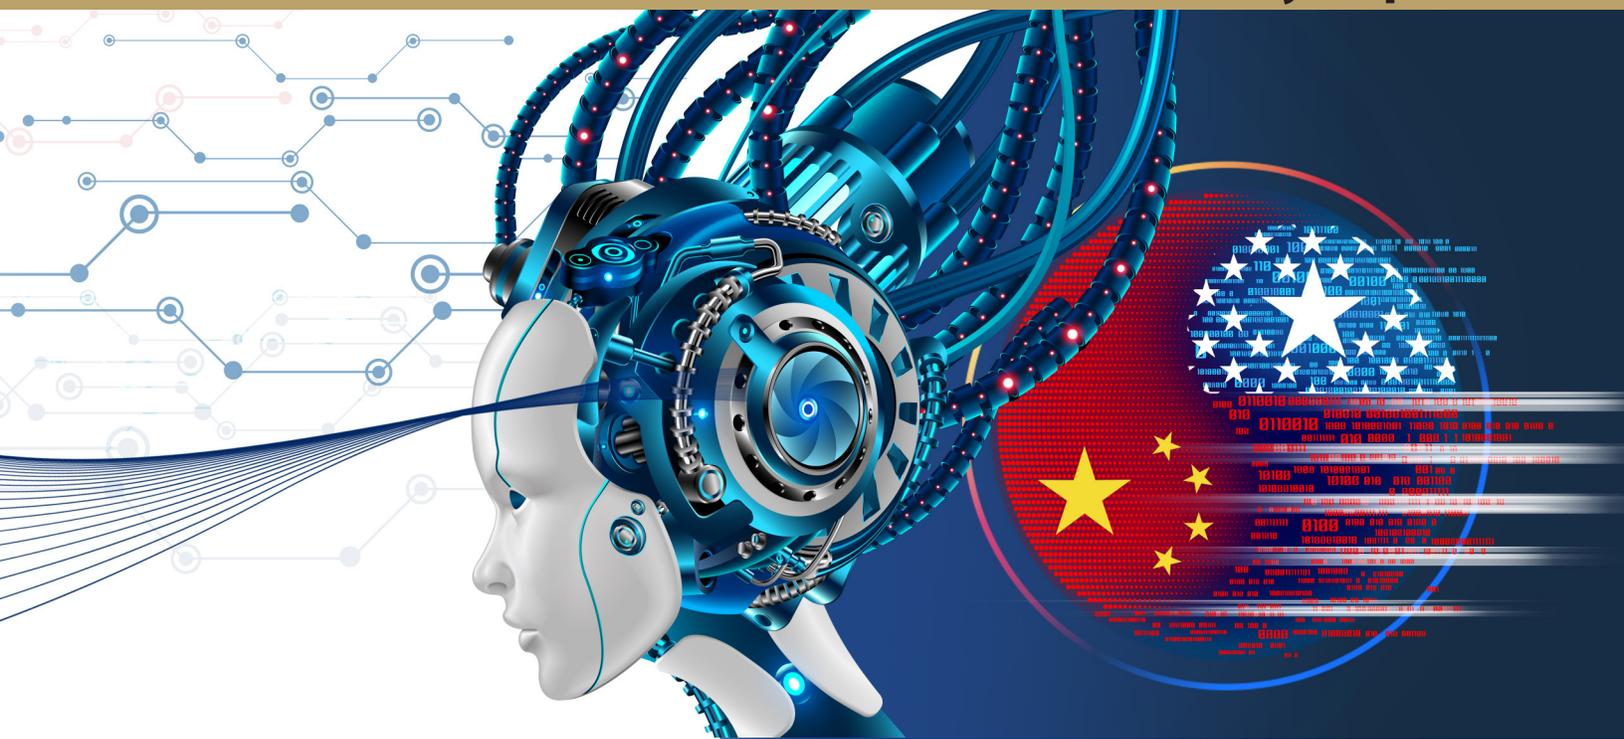


AN ENTWINED AI FUTURE RESISTANCE IS FUTILE

National Security Report



Christine Fox

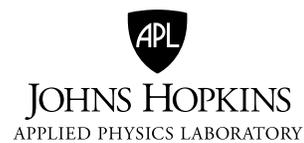


JOHNS HOPKINS
APPLIED PHYSICS LABORATORY

AN ENTWINED AI FUTURE

Resistance Is Futile

Christine Fox¹



¹ With contributions from Ralph Semmel (APL), Thayer Scott (APL consultant), Ashley Llorens (APL), John Piorkowski (APL), Lorand Laskai (Center for Security and Emerging Technology), Seth Weiner (APL), and Melissa Terlaje (APL).

Copyright © 2020 The Johns Hopkins University Applied Physics Laboratory LLC. All Rights Reserved.

The views in this document reflect the opinions of the author alone and do not represent any institutional position held by APL.

Contents

Foreword	v
Summary	vii
An AI Race?	1
How to Protect?	2
Research and Human Capital.....	2
Infrastructure	5
Data	6
Applications	8
Testing for Trust	9
National Security Opportunities	10
Conclusion	11
Bibliography	13
About the Author	19

Foreword

This paper is part of the “Measure Twice, Cut Once: Assessing Some China–US Technology Connections” research series sponsored by the Johns Hopkins University Applied Physics Laboratory.

As competition has intensified between the United States and China, actions to disengage their technology establishments from one another have also intensified. The two countries’ systems for research and development, production, and sale of cutting-edge technologies have been substantially, though by no means uniformly, commingled. More recently, there have been concerted efforts by both nations’ governments to reverse some or all of that commingling. Policymakers’ priorities include perceived risks to national security, worry about economic disadvantage from proliferation, and concern about uses of technologies that intentionally or indifferently may harm civil liberties or the environment.

To explore the advisability and potential consequences of decoupling, the Johns Hopkins University Applied Physics Laboratory commissioned papers from experts in specific technology areas. In each of these areas, the authors have explored the feasibility and desirability of increased technological separation and offered their thoughts on a possible path forward. Other papers in this series include:

- *Two Worlds, Two Bioeconomies: The Impacts of Decoupling US–China Trade and Technology Transfer* by Rob Carlson and Rik Wehbring
- *The History and Future of US–China Competition and Cooperation in Space* by Matthew Daniels
- *Symbiosis and Strife: Where Is the Sino–American Relationship Bound? An Introduction to the APL Series “Measure Twice, Cut Once: Assessing Some China–US Technology Connections”* by Richard Danzig and Lorand Laskai
- *Cutting off Our Nose to Spite Our Face: US Policy toward Huawei and China in Key Semiconductor Industry Inputs, Capital Equipment, and Electronic Design Automation Tools* by Douglas B. Fuller
- *The Telecommunications Industry in US–China Context: Evolving toward Near-Complete Bifurcation* by Paul Triolo
- *Addressing the China Challenge for American Universities* by Rory Truex
- *US–China STEM Talent “Decoupling”: Background, Policy, and Impact* by Remco Zwetsloot

Summary

Both the importance of artificial intelligence (AI) and the difficulty of controlling its dissemination derive from its character as a general-purpose technology. Although most of the foundational work on AI was initially pursued by the US government, the largest investments and developments of the past two decades have taken place in the commercial sector, with the results publicly available in open source for use by many thousands of programmers collaborating across national boundaries. AI by its nature cannot be stored in a warehouse. Though AI will ultimately, in some form, be present in many, perhaps most, weapons systems, it is not a weapons system. AI is becoming as pervasive and accessible as, for example, electricity or digital computing or even arithmetic in the last century.¹ As such, we do not think its fundamentals can or should be controlled by government action.

We do think, however, that some particular applications and data sets relevant to national security capabilities can and should be controlled. It is also sensible to limit Chinese access to some high-end semiconductor chips required for the most sophisticated AI applications and, especially, to retain US and allied dominance over the tools needed to manufacture those chips.

The preferred Cold War tools for blunting an adversary's technological advances—including classification and export restriction—are not as applicable to most forms of AI (though the Commerce Department is testing that proposition with recent export controls). Many key AI technologies are shared freely through open-source resources such as GitHub and Google's TensorFlow.² Beyond that, Chinese access to information, hardware, and software can occur by a variety of means beyond purchase from the United States. These include cyber theft, coercive joint ventures with non-Chinese companies, access to US universities and market programs, and transactions with third-party countries. Plugging these gaps will be of greater benefit than attempting to decouple wholesale from China's AI ecosystem with blunt policy tools.

We believe that Chinese advantages—large and robust markets, talented people, and skilled training programs—will inevitably make China a strong AI competitor. The United States must invest comparably in the robustness of its markets, the talent of its population, and the strength of its training programs. It is more productive to make America stronger than it is to make China weaker.³ Attempting the latter by decoupling weakens America by isolating us from much of the world that will continue to do business with China while cutting off America's access to a major source of the AI talent and innovation.

This paper explores the US–China AI relationship in its various dimensions and latest developments; then considers what a productive and mutually beneficial relationship would be; and, in closing, offers thoughts on how to reconcile these goals, where possible, with our most critical national security imperatives.

¹ AI has been characterized as a “dual-use” technology but, as the National Security Commission on Artificial Intelligence *Interim Report* points out, that term does not really capture what will be the ubiquity of AI in all spheres of life.

² Dean and Monga, “TensorFlow.”

³ Major theme of Danzig et al., *Preface to Strategy*.

An AI Race?

It is tempting to see AI through the “race” paradigm, echoing the twentieth century competitions in nuclear weapons and space exploration. But, in this case, there is no finish line, the course is ill defined, and the runners’ track lanes are blurry and overlapping. American and Chinese academic research communities routinely engage and publish with each other—along with other international partners. There is no AI equivalent of the Apollo lunar landing to strive toward. Our competition with China on AI is less a race—a 100-yard dash—than a track and field competition made up of different events. Inputs to those events would include research capabilities, access to talent, time to market, and, from a national security perspective, speed of adoption and deployment.

The United States and the People’s Republic of China (PRC) have different advantages in different aspects of this competition.¹ Kai Fu Lee, a leading AI scientist and investor based in Beijing, is famously bullish on China’s AI performance and potential. Lee’s basic premise is that Chinese access to a larger share of data and more aggressive commercialization are likely ultimately to outweigh the US advantages in basic research and innovation.² Eric Schmidt, former Google CEO and co-chair of the National Security Artificial Intelligence Commission (NSAIC), cited “integration innovation”—adapting or commercializing what had been created elsewhere—as a key Chinese advantage.³ So too are China’s widespread embrace of STEM education and its ability through industrial policy to mobilize and sustain long-term government-directed funding, separate from market demands.⁴

¹ Imbrie, Kania, and Laskai, *Question of Comparative Advantage*.

² Parker, “Battle for Supremacy.”

³ Smith, “AI for National Security.”

⁴ For an examination of China’s government AI policies, see Ding, *Deciphering China’s AI Dream*.

As a foundational technology, AI’s general promise should not be locked up by classification and commerce restrictions.

However, the United States has distinct competitive advantages and strengths of its own: our values; system of government; human capital educated in superb universities; traditionally welcoming environment for immigrants; economic engine and vibrant commercial tech enterprise; favorable position in the international order and military alliances with the most technologically advanced nations; and innovative national security research labs and companies.⁵ Top Chinese experts acknowledge that many of these advantages have translated into an early and, so far, enduring lead in AI, especially with regard to foundational research, semiconductors, and AI frameworks.⁶ Scholars at Georgetown’s Center for Security and Emerging Technology (CSET) have shown how AI-related innovations, such as image generation, complex strategy games, and language understanding/generation, all come from labs focused on fundamental research and development (R&D) located primarily in the United States—not from profit-hungry start-ups in China.⁷ Despite the recent focus on Chinese implementation successes, it remains true that fundamental R&D remains a strategic US advantage. One way of viewing the present challenge is to ask whether America can continue to engage China in productive ways without ceding America’s advantages in these areas.

⁵ Danzig et al., *Preface to Strategy*.

⁶ Hickert and Ding, “Read What Top Chinese Officials Are Hearing.”

⁷ Laskai and Toner, “Can China Grow Its Own AI Tech Base?”

How to Protect?

During the Cold War, the US development of stealth radar-evading designs and materials remained a state secret, and the United States maintained its monopoly on stealth technology for nearly two decades, despite the many thousands of individuals in and out of government working on those technologies.⁸ But the closely held R&D projects of the Cold War bear little resemblance to the generally applicable (and widely available) AI technologies of today.⁹ AI is a foundational technology and, as a foundational technology, its general promise should not be locked up by classification and commerce restrictions.

However, as its overall relationship with China has deteriorated, the United States has increasingly turned to these traditional tools to protect AI. The Export Control Reform Act (ECRA) passed in August 2018 requires the Commerce Department to consider controlling “emerging” and “foundational” technologies—not specified in the law—that are “essential to the national security of the United States.”¹⁰ A follow-up Commerce Department notice soliciting public feedback listed AI among a group of such technologies, and later, AI-based geospatial analysis applications were controlled for the first time.¹¹ However, the most recent, and sweeping, set of Commerce controls issued in April 2020 notably—and appropriately—did not

include AI.¹² We believe that attempting to impose, for national security purposes, traditional export controls on a general-purpose technology like AI is neither feasible nor effective. Conversely, AI-related hardware, software, and data sets with military or intelligence applications are more legitimate candidates for restrictions. For the most part, concerns regarding the export of technologies with military or intelligence applications are already addressed by existing US regulations governing end users and purposes (as opposed to general AI technology), but some gaps and vulnerabilities remain.¹³

Research and Human Capital

Top talent is crucial to making advances in AI and is in short supply.¹⁴ Xi Jinping has called talent “the first resource” of China’s innovation push.¹⁵ China’s 2017 New-Generation AI Development Plan (AIDP) declared that developing high-end talent is “of the utmost importance” to China’s AI development.¹⁶ The country’s leadership has, accordingly, invested heavily in increasing its AI talent pipeline in recent years, an effort that appears to be achieving results.¹⁷ Seeking to tap this growing supply of AI talent, a number of American tech giants have established research centers in China.

⁸ Tirpak, “Two Decades of Stealth.”

⁹ Horowitz, “Artificial Intelligence.”

¹⁰ 50 USC § 4817. The ECRA’s emerging and foundational technologies provision originally stemmed from calls to fill a perceived regulatory gap between the Commerce Department’s controls on technology transfer to foreign persons and CFIUS (Committee on Foreign Investment in the United States) controls on foreign investments in US companies. For a discussion of the regulatory gap, see Brown and Singh, *China’s Technology Transfer Strategy*.

¹¹ Department of Commerce, “Review of Controls”; and Department of Commerce, “Addition of Software.”

¹² Behsudi, “A Potential Game-Changer”; and Department of Commerce, “Expansion of Export.”

¹³ Flynn, *Recommendations on Export Controls*; Leiter, Gerkin, and Klein, *Commerce Department*; and Aitel, “We Need a Drastic Rethink.” To illustrate, the federal regulation adding machine learning-enabled geospatial imagery applications to the Commerce Control List relied on long-standing authorities already available to control the export of dual-use software applications. Nevertheless, critics argue that the new restriction is both overly broad and probably ineffective. While ample regulatory tools exist, the challenge for policymakers lies in appropriately and effectively applying these tools to achieve their stated goals without causing unnecessary harm.

¹⁴ Metz, “Tech Giants.”

¹⁵ Zwetsloot and Peterson, “The US-China Tech Wars.”

¹⁶ Webster et al., “Full Translation.”

¹⁷ Dantong Ma, “China’s AI Talent Base.”

When Google established a center in Beijing, its chief AI scientist pointed out that the winning teams of the ImageNet Challenge—a competition that evaluates algorithms for object detection and image classification—had been largely composed of Chinese researchers.¹⁸

Drawing on the world's best and brightest and pursuing international engagement in AI development and markets ultimately play to America's strengths.

However, even as China produces a growing workforce for AI, the most capable of those individuals are choosing to attend school and then work in the United States. For an AI researcher, the United States is a very attractive place to be for many reasons, not the least of which is that the United States dominates the leading AI conferences, the most common AI frameworks, and leading research professional societies. MacroPolo found about three-quarters of all Chinese authors that have presented at NeurIPS, one of the top conferences for AI research, are working outside China, mainly in the United States.¹⁹ Reflecting on the AI brain drain, a top expert at the Chinese Academy of Sciences, Tan Tieniu, warned that a shortage of mid- and high-end talent is the main bottleneck in China's AI development.²⁰

For their part, Chinese companies have established similar AI research centers in the United States. For example, Alibaba and Tencent employ researchers in both Silicon Valley and Seattle. SenseTime, a Chinese AI company doing computer vision and deep learning, has an AI-based health lab in New Jersey, collaborates with MIT on machine

intelligence, and engages in “synchronous development” between its China and Silicon Valley locations.²¹ (SenseTime was one of the Chinese companies later sanctioned by the United States for aiding human rights violations²²; MIT is currently reviewing the relationship.)

Drawing on the world's best and brightest and pursuing international engagement in AI development and markets ultimately play to America's strengths. Such efforts strengthen the R&D community, and they strengthen US national security providers in the long-term struggle with our major competitors. However, with respect to China, the United States must account for the fact that China does not necessarily acknowledge the distinct separation between the private, public, and nonprofit sectors that defines Western societies. Accordingly, the United States should assume that AI knowledge obtained by any Chinese national, company, or organization is probably also accessible to the Chinese government for intelligence, domestic surveillance, or military purposes.²³

Shared basic research in AI goes beyond the investments of commercial tech companies in cross-nation research centers. Currently, more than 160,000 Chinese study science, technology, engineering, or math (STEM) subjects at US universities²⁴—a number that has roughly doubled over the last decade.²⁵ US academic institutions,

²¹ Hannas and Chang, *China's Access to Foreign AI Technology*.

²² Department of Commerce, “Addition of Certain Entities.”

²³ For example, China's National Intelligence Law of 2017 states that “any organization or citizen shall support, assist, and cooperate with state intelligence work according to law.” Moreover, “state intelligence work organs, when legally carrying forth intelligence work, may demand that concerned organs, organizations, or citizens provide needed support, assistance, and cooperation.” The scope of this obligation is potentially very broad because key terms like “intelligence work” are not defined by the law. Tanner, “Beijing's New National Intelligence Law.”

²⁴ Granovskiy and Wilson, *Foreign STEM Students*.

²⁵ Statista, “Number of College and University Students.”

¹⁸ Li, “Opening the Google AI China Center.”

¹⁹ Dantong Ma, “China's AI Talent Base.”

²⁰ Hickert and Ding, “Read What Top Chinese Officials Are Hearing.”

under near-constant financial pressure, are eager to accept international STEM students who typically pay full tuition and provide a steady supply of researchers and teaching assistants.²⁶

The United States cannot cut itself off from some of the best minds working in this area. Many of those minds are Chinese.

There are liabilities associated with relying so heavily on Chinese nationals for revenue and research. In China, the relationships between academic institutions and the government are intertwined in ways that are inconceivable in America. The notion of academic freedom, including freedom from government interference, differs greatly from Chinese norms—such differences are generally not appreciated by students from either country.²⁷ A 2019 report commissioned by the National Science Foundation on fundamental research security observed that China’s activities to gain access to US academic research involve a mix of “reward, deception, coercion, and theft.”²⁸ China has more than 200 talent-recruitment programs, the most notable of which is the Thousand Talents Program. This program seeks to lure academics engaged in important scientific research in topics such as AI with offers of high salaries, research funding, and support facilities.²⁹

However, even studies that expose and warn of these Chinese practices do not advocate for the United States to restrict access of Chinese students to our universities. The intellectual benefit of international collaboration in important new areas

like AI is too great to ignore.³⁰ Surveys show that up to 90 percent of Chinese students who earn their PhDs stay in the United States for at least five years, contributing to our economy while depriving China of their brainpower.³¹ This is a gain for the United States and a loss for China as many of China’s best and brightest choose to work and live in the United States and not China. However, as the US government gradually restricts the flow of Chinese graduate students into the United States, Chinese officials see an opportunity to reverse this talent outflow and retain more of the top AI talent in China.³²

Nonetheless, academic funding organizations and US study participants should be informed more clearly about Chinese practices and incentivized, or directed, to conduct certain kinds of AI research in more secure ways. Many universities, for example, are very aware of the risks and have put in place more stringent controls.³³ Individual professors, labs, and university departments have established reputations with government funding organizations and are known and trusted. We believe that a more systematic approach is needed, standardized by government direction if required, to establish these and other best practices across academia. Nonetheless, the United States cannot cut itself off from some of the best minds working in this area. Many of those minds are Chinese.

²⁶ Makala Skinner’s article (“The Financial Risk of Overreliance on Chinese Student Enrollment”) highlights the risks to university finances of relying so heavily on Chinese students.

²⁷ JASON, *Fundamental Research Security*.

²⁸ JASON, *Fundamental Research Security*, 21.

²⁹ Portman and Carper, *Threats to the U.S. Research Enterprise*.

³⁰ “There is a long and illustrious history of foreign-born scientists and engineers training and working in the United States, and they make essential contributions to our preeminence in science, engineering and technology today. Maintaining that leading position will require that the United States continues to attract and retain the best science talent globally” (JASON, *Fundamental Research Security*, 2).

³¹ Zwetsloot, Feldgoise, and Dunham, *Trends in U.S. Intention-to-Stay Rates*.

³² Qian and Hualing, “News Analysis.”

³³ MIT, for example (MIT Office of the Vice President for Research, “Export Control”).

Infrastructure

AI requires specialized chips that are powerful, efficient, and optimized for advanced machine learning algorithms. These “AI chips” are forecast to comprise up to 20 percent of the \$450 billion semiconductor chip market by 2025.³⁴ Currently less than a fifth of the semiconductors used in China are made domestically.³⁵ Major US semiconductor companies generate a disproportionate portion of profits from the Chinese market—and much of those profits are directed toward R&D that sustains America’s leading edge in AI. The market for semiconductor manufacturing equipment and tooling—the more strategically critical “choke point” capability—is even more lopsided, dominated by the United States, Japan, and the Netherlands.³⁶

The Chinese government is keenly aware of these imbalances and wants to change them. Officials understand that cutting-edge AI relies on harnessing large amounts of computational power, and without chip independence, the country’s AI ambitions, in the words of Wang Yu, a professor at Tsinghua University, would be the equivalent of “building a house on a foundation of sand.”³⁷ As a government white paper succinctly put it, “No chip, no AI.”³⁸

Given China’s reliance on foreign suppliers, the Chinese government fears that the United States will cut off its supply of AI chips. Two years ago, in 2018, the United States imposed a supply ban on ZTE and, more recently, in 2019, banned several Chinese AI start-ups because of their roles in human rights violations.³⁹ And in May 2020, the Commerce Department introduced a rule change

that will block companies from around the world from using American-made machinery and software to produce or design chips for Huawei or its entities.⁴⁰

These kinds of sanctions—viewed by China as insulting as well as economically harmful—have deepened the resolve of Chinese officials to push even harder to accelerate their chip development industry.⁴¹ Chinese industry leaders, too, have joined the call for prioritizing in-house “core technologies.” After banning ZTE, Alibaba CEO Jack Ma declared, “Big enterprises have an important responsibility. If we do not master the core technologies, we will be building roofs on other people’s walls and planting vegetables in other people’s yards.”⁴² Since then, the company has invested heavily in developing chips for AI applications.⁴³ As the United States targets a growing number of Chinese companies’ access to critical US semiconductors, those companies understand their future prospects will be increasingly entwined with the government’s push for chip independence.⁴⁴

An AI development community dependent on Chinese chips would represent a considerable American economic and security risk. China is nothing if not opportunistic.

China faces significant hurdles when it comes to matching the chip development capabilities of the United States and our allies. As recently as last year, Reuters reported that industry insiders characterized Chinese chip companies as “relatively

³⁴ MacroPolo, “Big Picture.”

³⁵ Horwitz and Jiang, “China Chip Industry Insiders.”

³⁶ Fuller, *Cutting off Our Nose*.

³⁷ Science and Technology Daily, “In Developing AI Chip.”

³⁸ You and Shaojun, “AI Chip Technologies.”

³⁹ Li, “China ‘Strongly Urges’ US.”

⁴⁰ Swanson, “U.S. Delivers Another Blow.”

⁴¹ Simons, “China Tipped to Accelerate.”

⁴² “Alibaba’s Jack Ma,” *Shanxi Evening News*, quoted in Segal, “Seizing Core Technologies.”

⁴³ Knight, “Alibaba’s ‘Honey Badger.’”

⁴⁴ Laskai, “Why Blacklisting Huawei Could Backfire.”

backward,' lacking in talent and 'requiring a long time to catch up.'⁴⁵

This situation is addressed in depth in Douglas Fuller's "Measure Twice, Cut Once" companion paper on the semiconductor industry.⁴⁶ We note our general conclusion is that China is unlikely to catch up quickly in semiconductor manufacturing, but it is making progress. For more basic chip technology, such as memory chips or flash chips, China already has competitive capabilities. Our efforts to gain leverage on China by turning down the spigot of chips is incentivizing the Chinese to put even more investment in and focus on developing a capable indigenous chip industry. If they are successful, we would face more than just economic competition. By gaining control of the supply chain, or even capturing a large part of it, China could infuse the chips with means to exfiltrate information from any resident device.⁴⁷ An AI development community dependent on Chinese chips would represent a considerable American economic and security risk. China is nothing if not opportunistic. However, we should recognize that our short-term security gains may lead to long-term consequences. Rather than denying China access to US chips, the United States should carefully control its chip manufacturing knowledge, talent, and equipment. By giving China reliable and mutually beneficial access to US chips, we can reduce their incentive to develop their own capability, protecting both the US market and supply chain in the long run.

Data

Data is essential to current machine learning techniques. Some say that when it comes to AI, data is the new oil and computing power the new combustion engine. Although future AI may rely less on data, currently AI systems are data hungry

and will probably remain so for the foreseeable future. There is no question that China has access to enormous quantities of data. The size of the Chinese population, coupled with a willingness to extract and use personal and private information, gives PRC companies extremely large data sets to work with.⁴⁸ While Chinese and US citizens share private information easily and willingly with private companies, Chinese urban residents conduct much more of their personal, commercial, and professional transactions through smartphones as compared with the United States.⁴⁹ And in China, the government also has access to private company data.

Unlike most algorithms, data can be protected and controlled. In the national security realm, it should be.

This Chinese data advantage—quantity and access—has its limits. Some data is more useful and usable than others depending on its quality, depth, and diversity.⁵⁰ China's government and industry have enormous insight into the browsing and spending habits of the population, but these data are of limited utility in other contexts.⁵¹ Even with vast quantities of data, effectively using that data can be challenging. Robin Li, founder and CEO of Baidu, said the biggest challenge facing new technologies like AI is "data islands." "We have a huge amount of data, but this data is often fragmented," he said last year.⁵² Chinese government experts also worry that stovepiping and bureaucratic competition in the Chinese government are limiting the country's

⁴⁵ Horwitz and Jiang, "China Chip Industry Insiders."

⁴⁶ Fuller, *Cutting off Our Nose*.

⁴⁷ See, for example, Robertson and Riley, "The Big Hack."

⁴⁸ Lee, *AI Superpowers*.

⁴⁹ Sheehan, "Much Ado about Data"; see also Ding, *Deciphering China's AI Dream*.

⁵⁰ Sheehan, "Much Ado about Data."

⁵¹ Zwetsloot, Toner, and Ding, "Beyond the AI Arms Race."

⁵² Xinxin and Weiwei, "Commissioner Li Yanhong."

ability to effectively share and utilize data sets.⁵³ While the Chinese government has vast amounts of data—over 80 percent of all China’s data, according to China’s Premier Li Keqiang—the structure and secretive nature of the Chinese government might make it difficult to effectively leverage that data.⁵⁴

Overall, the US data ecosystem is more open and thus more effective at leveraging quality data.⁵⁵ US companies and government agencies have invested in enterprise software and digitization to standardize data sharing, something that has not occurred in China.⁵⁶ Last year, Congress passed the OPEN Government Data Act, which requires government agencies to make nonsensitive data available in machine-readable formats.⁵⁷ Chinese experts have urged the Chinese government to pass its own version of the OPEN Government Data Act, though no action has been forthcoming.⁵⁸

Unlike most algorithms, data can be protected and controlled. In the national security realm, it should be. Intelligence data, performance data, and other types of data unique to national security are already commonly classified. Importantly, national security applications often require narrower and more specialized data sets than the “big data” used in the private sector. For example, there are not many high-quality images of explosive mines on the ocean floor or missile launchers in the woods as compared with images of pets or dogs or people. These small data sets require different AI algorithms that leverage specialized domain knowledge to “fill in the gaps” created by sparse or missing data. There will be cases in which it is critically important to

protect not just national security data sets but also these specialized variants of AI algorithms.

A military force that can immediately recognize that its data and associated algorithms have been poisoned, ideally while also providing a “cure,” could have a significant advantage in future conflicts.

To trust AI-enabled capabilities, it is also essential to protect the integrity of the data. Some national security applications rely on the use of large, internationally shared data sets. While providing a valuable resource for bootstrapping development, these open data sets can also introduce hard-to-identify vulnerabilities into AI algorithms. The Intelligence Advanced Research Projects Activity (IARPA) “Trojans in Artificial Intelligence” (TrojAI) project—recalling the Trojan Horse of Greek mythology—studies many of these vulnerabilities, including data poisoning. A well-publicized IARPA study looked at how AI in an autonomous vehicle could learn to read and react to different kinds of traffic signs. A bad actor—criminal, national, or otherwise—could insert images of some altered stop signs into the data set and relabel them as speed limits. This subtle change could cause a self-driving car to drive right through a stop sign.⁵⁹ On a larger scale, the result might be carnage at busy intersections or military convoys directed on trips to nowhere (or worse). One could easily conceive of data “backdoors” designed to change the classification of an object. At the Johns Hopkins University Applied Physics Laboratory (APL), our researchers developed a small, physical “patch” that, if worn, causes a classifying AI algorithm to identify a human as a teddy bear. These capabilities could deter an adversary from relying on their

⁵³ Yu, “Difficulties of Data Sharing.”

⁵⁴ Da and Baoguo, “Viewing US Thinking.”

⁵⁵ According to the McKinsey Global Institute, the Chinese government ranks ninety-third in the world for data openness (the United States ranks eighth) (Barton et al., “Artificial Intelligence”).

⁵⁶ Sheehan, “Much Ado about Data.”

⁵⁷ Chappellet-Lanier, “OPEN Government Data Act.”

⁵⁸ Da and Baoguo, “Viewing US Thinking.”

⁵⁹ IARPA, “Trojans in Artificial Intelligence.”

AI-enabled capabilities—or deter US forces from relying on theirs.

A growing number of research efforts aim at enhancing the ability to identify evidence of data poisoning and other potential failure modes in AI algorithms. The fruits of these research efforts should be carefully guarded. A military force that can immediately recognize that its data and associated algorithms have been poisoned, ideally while also providing a “cure,” could have a significant advantage in future conflicts.

Data are essential to AI research, development, and applications. The quality and integrity of the data underpin the value of the AI application. The United States needs to work with international and commercial partners to establish international standards governing the sharing of large data sets across national boundaries and industries. Through stronger collaboration, the United States can better manage the pedigree of the data underlying AI applications, whether they are used for civic, commercial, or national security purposes.

But even with standards and international oversight, we should recognize that bad data will be an issue for all AI applications. In the national security realm, purposeful data poisoning poses great risks to AI-enabled capabilities. All attempts should be made to protect algorithms that can identify and protect against poisoned data.

Applications

AI is an enabling capability. A very powerful one, but only an enabler. Without an application, AI remains sets of numbers and equations interesting to computer scientists but not many others. Accordingly, the competition for superiority in AI is likely to be a battle over applications. Domain expertise—how well each country can apply AI technology to improve a given field of society, governance, or security—will be a major factor.

Over time, China’s well-coordinated and aggressive advocacy for international standards that reflect its interests and values will bear fruit at America’s expense.

Disconnecting the United States from the global ecosystem that develops AI applications—an ecosystem from which it is impossible to exclude China—would ultimately hurt the interests and well-being of our own citizens. For example, AI applications in the health field have the potential to radically improve health care for everyone. Today, data analytics and AI are being used to do everything from refining patient populations for targeted treatment, to assessing X-rays, to automating care delivery. In another application, “smart cities” use AI-enabled automation to manage traffic, schedule mass transit, and ensure services are delivered to their populations. These cities are expected to be more energy efficient and have less adverse impact on the environment.

Without constant engagement with international standards-setting organizations and collaboration with other nations, the United States will not be able to influence global norms on these and other applications. Otherwise obscure international groups can “bake in” technical specifications that may last for decades. China has set an explicit goal of becoming “a standards-issuing country” and coordinates standards work across government, industry, and academia.⁶⁰ The Chinese government holds pre-meetings with industry (expected to be advocates as national champions), sends large delegations, and creates voting blocs with Belt and Road Initiative partners. In March 2020, a letter signed by seventeen US senators from both political parties voiced concern over China’s use of

⁶⁰ Gorman, “U.S. Needs to Get in the Standards Game.”

international bodies to enshrine its preferred norms and rules for advanced surveillance technology.⁶¹

Over time, China's well-coordinated and aggressive advocacy for international standards that reflect its interests and values will bear fruit at America's expense. The US approach to standardization has been bottom-up, stakeholder driven, and generally resistant to central planning—a posture that is no longer sufficient.⁶² Withdrawing from exchange and collaboration with China effectively means withdrawing from dealings with major swaths of the world, including increasingly traditional US trading and security partners. Not only would we be potentially ceding a good portion of the market, we would be ceding the argument about the appropriate role of AI technology in society.

There is a common interest in ensuring that AI on our roads and in our homes is trustworthy and performs as expected.

Testing for Trust

In addition to standards for data and applications, standards for how to test an AI-enabled capability are equally important and less understood. Use of any of these AI-enabled capabilities—military, commercial, or otherwise—requires trust that they will perform as expected. As a general proposition, it is difficult to predict definitively how an AI algorithm will perform. We train it on data or, for more sophisticated algorithms, give it goals and let it run, but we do not know exactly how it will decide what it will do. To have confidence that it will perform as we intend and produce the outcomes we seek, we will need to test it. But how can we ever be sure? The algorithm will continue to learn, and

performance will continue to change. These are fundamental questions that the AI development community is struggling with today.

Assuming that the world's two biggest economies will remain intertwined, invariably American companies and consumers will be using Chinese-developed AI at some point and vice versa. There is a common interest in ensuring that AI on our roads and in our homes is trustworthy and performs as expected. For us to be confident that a Chinese AI-enabled capability is safe to use, we need to understand how it was tested. Even better, we will engage with Chinese researchers to design testing approaches and establish testing norms that govern the products we adopt.

We will also need to have confidence in how warfighting capabilities were tested before using them—or facing them—on the battlefield. It is one thing to face a thinking adversary in war—it is another to face an unpredictable machine. In the brinkmanship that often accompanies conflict, it is important that both sides be confident in their ability to control their weapons. So, the key to trusted AI is testing.

At the same time, it is inevitable in conflict that both sides will attempt to undermine the capabilities of the other through any means possible. All nations that develop missile systems test them to ensure they will work as designed and then face adversaries who have developed countermeasures to degrade that performance. The same is true in cyber. We develop—and counter—each other's cyber capabilities, creating uncertainty in the reliability of our systems. We will do the same with AI.⁶³

Additionally, because AI is not entirely explainable, it engenders fear tied not to any adversary actions but to the nature of the technology itself. The key to overcoming that fear is trust, and the key to gaining trust is testing. We should share research practices

⁶¹ Portman et al. to Pompeo, March 11, 2020.

⁶² Gorman, "U.S. Needs to Get in the Standards Game."

⁶³ An ongoing Defense Advanced Research Projects Agency (DARPA) priority; Turek, "Explainable Artificial Intelligence."

that will enable both sides to create reliable AI capabilities. But we should also be prepared for the inevitable adversarial AI game where we undermine each other's capabilities. By sharing testing protocols, we run the risk of exposing our systems to adversary efforts to undermine them. There will be many AI-enabled technologies where we will have to protect the testing process and, certainly, the results. However, methodologies and performance standards can and should be shared to gain global trust in commercial products and weapons systems.

National security AI applications would be more developed with greater government funding. Commercial capabilities are essential but not sufficient. Domain expertise is necessary to translate those capabilities for use in the national security realm.

National Security Opportunities

Unlike those in health, commerce, energy, or infrastructure, the differences between “winning” and “losing” in the military arena are stark and potentially more consequential. AI can help decision-makers quickly gain access to precisely the right information. It can help push that information out to the tactical edge, enabling commanders to focus on the right things and avoid surprise. It can also free warfighting platforms from having to accommodate human limitations like g-strain in a fighter jet or passenger room and protection in a ground vehicle. AI-enabled capabilities won't get tired or heat-stressed or anxious, so they can be used longer and in harsher conditions. This will free those who master AI to develop military platforms with completely new designs.

With respect to relatively new technologies like AI, organizational capacity, or the ability and willingness to absorb a new approach, is as important as, if not more important than, financial capacity or funding.⁶⁴ Generally the US military, despite its reputation for bureaucracy and process, is less rigid and top-down as an organization than the People's Liberation Army. The Chinese military, however, is less wedded to conventional organizations and warfighting concepts. It has adopted new structures specifically to use AI and related technologies in asymmetric ways.⁶⁵

More recently, the Pentagon has begun to engage the commercial tech sector and bring new applications into the force. Some early integration of AI and related technologies has shown value to the defense enterprise, primarily in the areas of logistics, predictive maintenance, and intelligence, surveillance, and reconnaissance (ISR) analysis. Base algorithms and big data sets may be impossible to wall off from other nations. Various attributes of commercial systems must remain open, but if the Pentagon's applications have strategic value, it is reasonable for them to be closely held.

The need for these specialized AI capabilities reinforces the importance of renewing government investment in a field that has been largely dominated by the private sector in recent years. The government is generally in a better position to decide what to protect if it is engaged in its funding. The last cycle of major AI innovation was privately driven and publicly available. However, national security AI applications would be more developed with greater government funding. Commercial capabilities are essential but not sufficient. Domain expertise is necessary to translate those capabilities

⁶⁴ Described by Michael Horowitz as “adoption-capacity” theory (Ricks, “Michael Horowitz's Fine Study”).

⁶⁵ The People's Liberation Army set up a Strategic Support Force in December 2015 independent of other branches of the military for space, cyber, and electronic/information warfare (Costello and McReynolds, “China's Strategic Support Force”).

for use in the national security realm. Those activities require government funding.

We cannot be blind to the fact that China will take advantage of our open academic and commercial culture to exploit US accomplishments and, in some cases, turn them against our companies and, ultimately, our military forces.

Conclusion

America and China are certainly competing, and AI is a part, indeed a central part, of that competition. But AI knowledge, people, and commerce are continuously exchanged between the two powers. Americans and ultimately the entire global population benefit from AI advances propelled by this interchange. We need access to the best thinking on AI throughout the world, and much of that thinking is coming from talented Chinese academics and tech companies.

We believe that trying to disconnect R&D efforts from China will not propel America forward or significantly delay China's progress compared with ours. The only thing that is clear is that it will hold us both back. That said, we cannot be blind to the fact that China will take advantage of our open academic and commercial culture to exploit US accomplishments and, in some cases, turn them against our companies and, ultimately, our military forces. US companies and universities need more clarity on what to share and what to protect. Our government needs to insist that unclassified research at universities steer clear of sensitive work and follow the JASON report's recommendations for disclosure requirements and project assessments. The data and algorithms necessary to conduct complex military or intelligence operations

can and should be protected. However, the United States should continue to participate in the global AI community and, by doing so, take the opportunity to lead the world toward the constructive and ethical use of AI capabilities (and away from the authoritarian Chinese model), to include national security applications.

In the area of AI-related infrastructure, the United States continues to have significant advantages, none more so than in the market for semiconductor manufacturing equipment. We need to maintain our strategic edge in this advanced manufacturing arena. That requires protecting the highest-end technology from transfer or theft while also keeping the Chinese dependent on our products for as long as possible—a goal undermined by punitive sanctions and embargoes.

Finally, we return to where this essay began: AI is not a “thing,” a commodity that can be controlled. AI is composed of algorithms based on mathematics that will eventually be ubiquitous in everything we do. For the United States to be leaders in AI, we need to recognize its importance and invest in its development through education and training of our population and the robustness of our AI industry. By making these investments, collaborating intelligently in the international community, and rapidly adopting AI-enabled capabilities (particularly in national security), the United States will be better positioned to sustain our leadership in AI.

Bibliography

- Aitel, Dave. "We Need a Drastic Rethink on Export Controls for AI." *Net Politics* (blog). Council on Foreign Relations, January 21, 2020. <https://www.cfr.org/blog/we-need-drastic-rethink-export-controls-ai>.
- "Alibaba's Jack Ma on Developing Core Technologies Post-ZTE." [In Chinese.] *Shanxi Evening News*, April 24, 2018. <http://baijiahao.baidu.com/s?id=1598613211326939453&wfr=spider&for=pc>.
- Barton, Dominic, Jonathan Woetzel, Jeongmin Seong, Qinzhen Tian. "Artificial Intelligence: Implications for China." Paper presented at the 2017 China Development Forum. McKinsey Global Institute, April 2017. <https://www.mckinsey.com/~media/McKinsey/Featured%20Insights/China/Artificial%20intelligence%20Implications%20for%20China/MGI-Artificial-intelligence-implications-for-China.ashx>.
- Behsudi, Adam. "A Potential Game-Changer for China Export Controls." *Morning Trade* (Politico), April 28, 2020. <https://www.politico.com/newsletters/morning-trade/2020/04/28/a-potential-game-changer-for-china-export-controls-787183>.
- Brown, Michael, and Pavneet Singh. *China's Technology Transfer Strategy: How Chinese Investments in Emerging Technology Enable a Strategic Competitor to Access the Crown Jewels of U.S. Innovation*. Silicon Valley: Defense Innovation Unit Experimental, 2018. [https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_\(1\).pdf](https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_(1).pdf).
- Chappellet-Lanier, Tajha. "The OPEN Government Data Act Is Now Law." FedScoop, January 15, 2019. <https://www.fedscoop.com/open-government-data-act-law/>.
- Costello, John, and Joe McReynolds. "China's Strategic Support Force: A Force for a New Era." *China Strategic Perspectives*, edited by Phillip C. Saunders. Washington, DC: National Defense University Press, 2018. <https://ndupress.ndu.edu/Media/News/Article/1651760/chinas-strategic-support-force-a-force-for-a-new-era/>.
- Da, Shen, and Jia Baoguo. "Viewing US Thinking on Open Government Data from the Perspective of the Brand-New 'Open Government Data Act.'" Translated by Ben Murphy. *CAICT Official WeChat* (microblog), August 7, 2019 (translated August 9, 2019). https://docs.google.com/document/u/2/d/1_uEmssFvYCoSj0an6utFnIaG94_IBHjmXgUxi1Orvwk/edit.
- Dantong Ma, Joy. "China's AI Talent Base Is Growing, and Then Leaving." MacroPolo, July 30, 2019. <https://macropolo.org/chinas-ai-talent-base-is-growing-and-then-leaving/>.
- Danzig, Richard, John Allen, Phil DePoy, Lisa Disbrow, James Gosler, Avril Haines, Samuel Locklear III, James Miller, James Stavridis, Paul Stockton, and Robert Work. *A Preface to Strategy: The Foundations of American National Security*. National Security Perspective NSAD-R-18-038. Laurel, MD: Johns Hopkins University Applied Physics Laboratory, 2018. <https://www.jhuapl.edu/Content/documents/PrefaceToStrategy.pdf>.

- Dean, Jeff, and Rajat Monga. "TensorFlow - Google's Latest Machine Learning System, Open Sourced for Everyone." *Google AI Blog*, November 9, 2015. <https://ai.googleblog.com/2015/11/tensorflow-googles-latest-machine.html>.
- Department of Commerce, Bureau of Industry and Security. "Addition of Certain Entities to the Entity List." *Federal Register* 84, no. 196 (October 9, 2019). <https://www.federalregister.gov/documents/2019/10/09/2019-22210/addition-of-certain-entities-to-the-entity-list>.
- . "Addition of Software Specially Designed to Automate the Analysis of Geospatial Imagery to the Export Control Classification Number 0Y521 Series." *Federal Register* 85, no. 3 (January 6, 2020). <https://www.federalregister.gov/documents/2020/01/06/2019-27649/addition-of-software-specially-designed-to-automate-the-analysis-of-geospatial-imagery-to-the-export>.
- . "Expansion of Export, Reexport, and Transfer (in-Country) Controls for Military End Use or Military End Users in the People's Republic of China, Russia, or Venezuela." *Federal Register* 85, no. 82 (April 28, 2020). <https://www.federalregister.gov/documents/2020/04/28/2020-07241/expansion-of-export-reexport-and-transfer-in-country-controls-for-military-end-use-or-military-end>.
- . "Review of Controls for Certain Emerging Technologies." 15 C.F.R. 744. *Federal Register* 83, no. 223 (November 19, 2018). <https://www.federalregister.gov/documents/2018/11/19/2018-25221/review-of-controls-for-certain-emerging-technologies>.
- Ding, Jeffrey. *Deciphering China's AI Dream: The Context, Components, Capabilities, and Consequences of China's Strategy to Lead the World in AI*. Oxford, UK: Centre for the Governance of AI, Future of Humanity Institute, University of Oxford, 2018. https://www.fhi.ox.ac.uk/wp-content/uploads/Deciphering_Chinas_AI-Dream.pdf.
- Flynn, Carrick. *Recommendations on Export Controls for Artificial Intelligence*. Washington, DC: Center for Security and Emerging Technology, Georgetown University, February 2020. <https://cset.georgetown.edu/wp-content/uploads/Recommendations-on-Export-Controls-for-Artificial-Intelligence.pdf>.
- Fuller, Douglas B. *Cutting off Our Nose to Spite Our Face: US Policy toward Huawei and China in Key Semiconductor Industry Inputs, Capital Equipment, and Electronic Design Automation Tools*. National Security Report NSAD-R-20-059. Laurel, MD: Johns Hopkins University Applied Physics Laboratory, 2020.
- Gorman, Lindsay. "The U.S. Needs to Get in the Standards Game—With Like-Minded Democracies." *Lawfare* (blog), April 2, 2020. <https://www.lawfareblog.com/us-needs-get-standards-game-minded-democracies>.
- Granovskiy, Boris, and Jill H. Wilson. *Foreign STEM Students in the United States*. Report IF11347. Washington, DC: Congressional Research Service, November 1, 2019. <https://crsreports.congress.gov/product/pdf/IF/IF11347>.
- Hannas, Wm. C., and Huey-meei Chang. *China's Access to Foreign AI Technology: An Assessment*. Washington, DC: Center for Security and Emerging Technology, Georgetown University, September 2019. https://cset.georgetown.edu/wp-content/uploads/CSET_China_Access_To_Foreign_AI_Technology.pdf.

- Hickert, Cameron, and Jeffrey Ding. "Read What Top Chinese Officials Are Hearing about AI Competition and Policy." DigiChina project (blog). New America Cybersecurity Initiative, November 29, 2018. <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/read-what-top-chinese-officials-are-hearing-about-ai-competition-and-policy/>.
- Horowitz, Michael C. "Artificial Intelligence, International Competition, and the Balance of Power." *Texas National Security Review* 1, no. 3 (May 2018). <https://tnsr.org/2018/05/artificial-intelligence-international-competition-and-the-balance-of-power/>.
- Horwitz, Josh, and Sijia Jiang. "China Chip Industry Insiders Voice Caution on Catch-up Efforts." *Reuters*, June 13, 2019. <https://www.reuters.com/article/us-huawei-tech-usa-chip-catchup-analysis/china-chip-industry-insiders-voice-caution-on-catch-up-efforts-idUSKCN1TE1R4>.
- IARPA (Intelligence Advanced Research Projects Activity). "Trojans in Artificial Intelligence (TrojAI)." Accessed June 25, 2020. <https://www.iarpa.gov/index.php/research-programs/trojai>.
- Imbrie, Andrew, Elsa B. Kania, and Lorand Laskai. *The Question of Comparative Advantage in Artificial Intelligence: Enduring Strengths and Emerging Challenges for the United States*. Washington, DC: Center for Security and Emerging Technology, 2020. <https://cset.georgetown.edu/wp-content/uploads/CSET-The-Question-of-Comparative-Advantage-in-Artificial-Intelligence-1.pdf>.
- JASON. *Fundamental Research Security*. JSR-19-21. McLean, VA: The MITRE Corporation. December 2019. https://www.nsf.gov/news/special_reports/jasonsecurity/JSR-19-21FundamentalResearchSecurity_12062019FINAL.pdf.
- Knight, Will. "Alibaba's 'Honey Badger' AI Chip Company Hopes to Take on Bigger Beasts." *MIT Technology Review*. September 20, 2018. <https://www.technologyreview.com/2018/09/20/140083/alibas-honey-badger-ai-chip-company-hopes-to-take-on-bigger-beasts/>.
- Laskai, Lorand. "Why Blacklisting Huawei Could Backfire." *Foreign Affairs*, June 19, 2019. <https://www.foreignaffairs.com/articles/china/2019-06-19/why-blacklisting-huawei-could-backfire>.
- Laskai, Lorand, and Helen Toner. "Can China Grow Its Own AI Tech Base?" DigiChina project (blog). New America Cybersecurity Initiative, November 4, 2019. <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/can-china-grow-its-own-ai-tech-base/>.
- Lee, Kai-Fu. *AI Superpowers: China, Silicon Valley, and the New World Order*. New York: Houghton Mifflin Harcourt, 2018.
- Leiter, Michael E., Daniel J. Gerkin, and Nicholas A. Klein. *Commerce Department Will Move Forward with More Stringent Export Controls for Certain Emerging Technologies*. Washington, DC: Skadden, January 10, 2020. <https://www.skadden.com/insights/publications/2020/01/commerce-department-will-move-forward>.
- Li, Fei-Fei. "Opening the Google AI China Center." December 13, 2017. <https://www.blog.google/topics/google-asia/google-ai-china-center/>.

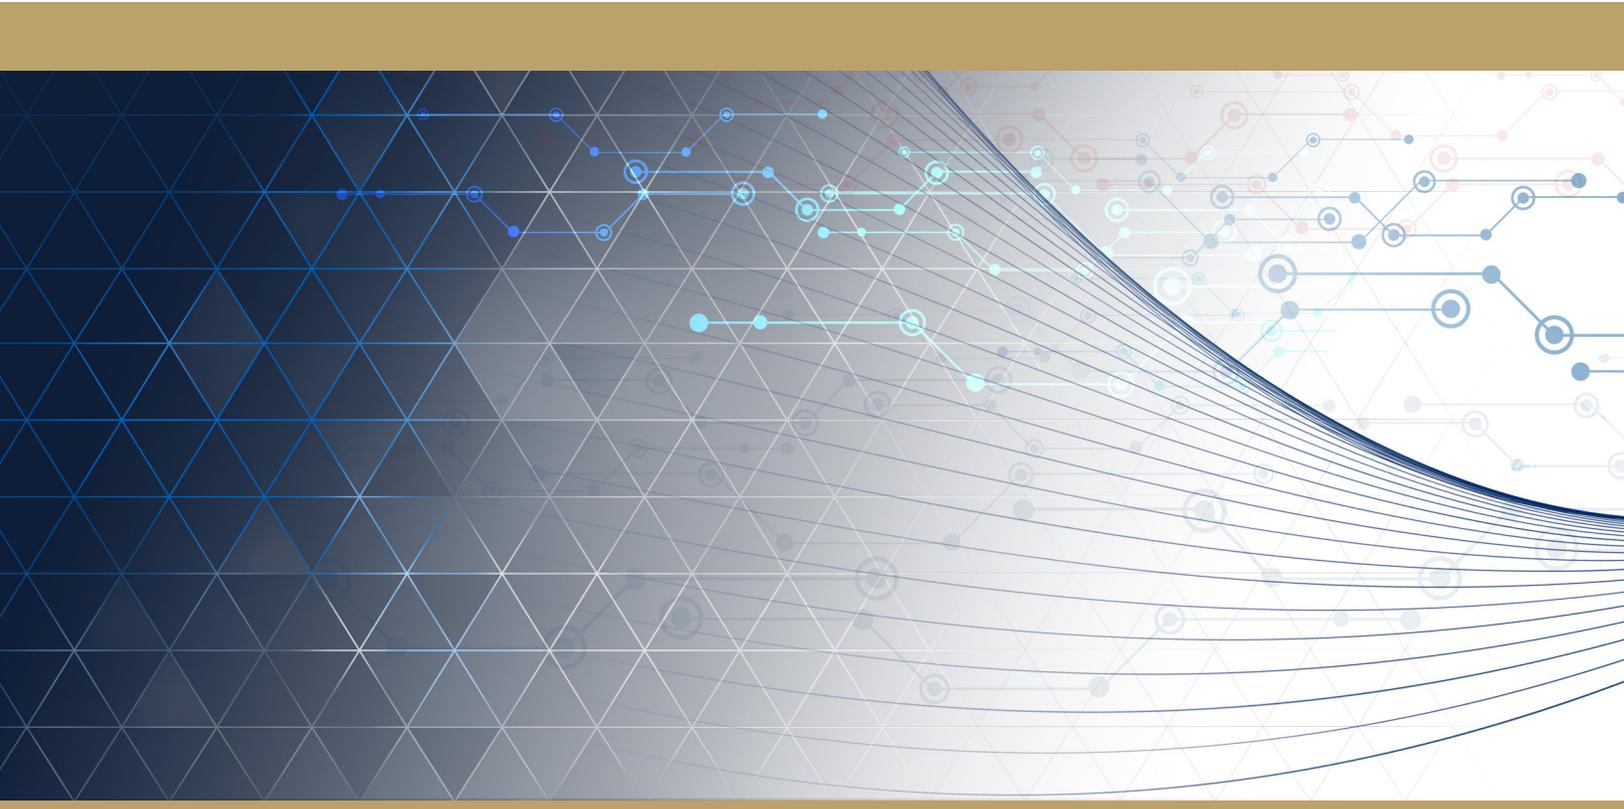
- Li, Yun. “China ‘Strongly Urges’ US to Remove Sanctions and Stop Accusing It of Human Rights Violations.” CNBC, October 8, 2019. <https://www.cnbc.com/2019/10/08/china-strongly-urges-us-to-remove-sanctions-and-stop-accusing-it-of-human-rights-violations.html>.
- MacroPolo. “Big Picture: AI Chips.” Accessed June 23, 2020. <https://macropolo.org/digital-projects/supply-chain/ai-chips/>.
- Metz, Cade. “Tech Giants Are Paying Huge Salaries for Scarce A.I. Talent.” *New York Times*, October 22, 2017. <https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html>.
- MIT Office of the Vice President for Research. “Export Control.” Accessed June 23, 2020. <https://research.mit.edu/integrity-and-compliance/export-control>.
- National Security Commission on Artificial Intelligence. *Interim Report*. November 2019. <https://drive.google.com/file/d/153OrxnuGEjsUvIxWsFYauslwNeCEkvUb/view>.
- Parker, Emily. “How Two AI Superpowers — the U.S. and China — Battle for Supremacy in the Field.” *Washington Post*, November 2, 2018. https://www.washingtonpost.com/outlook/in-the-race-for-supremacy-in-artificial-intelligence-its-us-innovation-vs-chinese-ambition/2018/11/02/013e0030-b08c-11e8-aed9-001309990777_story.html.
- Portman, Rob, Mark Warner, Richard Blumenthal, Tom Cotton, Christopher A. Coons, Cory Gardner, Steve Daines, et al. to Michael R. Pompeo. Letter, March 11, 2020. https://www.portman.senate.gov/sites/default/files/2020-03/China%20Warren%20AI%20Letter_0.pdf.
- Portman, Rob, and Tom Carper. *Threats to the U.S. Research Enterprise: China’s Talent Recruitment Plans*. Washington, DC: United States Senate, Permanent Subcommittee on Investigations, Committee on Homeland Security and Governmental Affairs. November 18, 2019. <https://www.hsgac.senate.gov/imo/media/doc/2019-11-18%20PSI%20Staff%20Report%20-%20China%27s%20Talent%20Recruitment%20Plans.pdf>.
- Qian, Peng, and Li Hualing. “News Analysis: Examining the Five Shortcomings of China’s AI Talent System.” Xinhua News Agency. August 28, 2019. http://www.gov.cn/xinwen/2019-08/28/content_5425310.htm. Translation available at <https://docs.google.com/document/d/1lJqEiettC0uHciGVVI94G0DHUZFm2IGZXIAfKt938uw/edit>.
- Requirements to Identify and Control the Export of Emerging and Foundational Technologies. 50 USC § 4817 (2018).
- Ricks, Thomas E. “Michael Horowitz’s Fine Study of How and Why Military Innovations Are Adopted.” *Foreign Policy*, July 25, 2014. <https://foreignpolicy.com/2014/07/25/michael-horowitzs-fine-study-of-how-and-why-military-innovations-are-adopted-2/>.
- Robertson, Jordan, and Michael Riley. “The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies.” *Bloomberg Businessweek*, October 4, 2018. <https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies>.

- Science and Technology Daily. "In Developing AI Chips, China Cannot Overemphasize Certain Parts over Others." [In Chinese.] July 8, 2019. http://tech.gmw.cn/2019-07/08/content_32979962.htm.
- Segal, Adam. "Seizing Core Technologies: China Responds to U.S. Technology Competition." *China Leadership Monitor*, June 1, 2019. <https://www.prclleader.org/segal-clm-60>.
- Sheehan, Matt. "Much Ado about Data: How America and China Stack Up." MacroPolo, July 16, 2019. <https://macropolo.org/ai-data-us-china/>.
- Simons, Hadlee. "China Tipped to Accelerate Domestic Chip Plans after ZTE Supply Ban." *Android Authority*, April 20, 2018. <https://www.androidauthority.com/china-domestic-chip-plans-zte-857392/>.
- Skinner, Makala. "The Financial Risk of Overreliance on Chinese Student Enrollment." World Education News + Reviews, December 17, 2019. <https://wenr.wes.org/2019/12/the-financial-risk-of-overreliance-on-chinese-student-enrollment>.
- Smith, Craig. "AI for National Security and the Challenge of China." *Forbes*, April 30, 2020. <https://www.forbes.com/sites/craigsmith/2020/04/30/ai-for-national-security-and-the-challenge-of-china/#12e2bf3d498a>.
- Statista. "Number of College and University Students from China in the United States from Academic Year 2008/09 to 2018/19." Accessed June 23, 2020. <https://www.statista.com/statistics/372900/number-of-chinese-students-that-study-in-the-us/>.
- Swanson, Ana. "U.S. Delivers Another Blow to Huawei with New Tech Restrictions." *New York Times*, May 15, 2020. <https://www.nytimes.com/2020/05/15/business/economy/commerce-department-huawei.html>.
- Tanner, Murray Scot. "Beijing's New National Intelligence Law: From Defense to Offense." *Lawfare* (blog), July 20, 2017. <https://www.lawfareblog.com/beijings-new-national-intelligence-law-defense-offense>.
- Tirpak, John A. "Two Decades of Stealth." *Air Force Magazine*, June 1, 2000. <https://www.airforcemag.com/article/0601stealth/>.
- Turek, Matt. "Explainable Artificial Intelligence (XAI)." Defense Advanced Research Projects Agency. Accessed June 25, 2020. <https://www.darpa.mil/program/explainable-artificial-intelligence>.
- You, Zheng, and Wei Shaojun, eds. "White Paper on AI Chip Technologies." Beijing Innovation Center for Future Chips, Tsinghua University, 2018. <https://www.080910t.com/downloads/AI%20Chip%202018%20EN.pdf>.
- Yu, Chen. "The Three Difficulties of Data Sharing: Unwilling, No Courage, Unlikely." [In Chinese.] *Science and Technology Daily*, October 29, 2019. http://digitalpaper.stdaily.com/http_www.kjrb.com/kjrb/html/2019-10/29/content_433625.html.
- Webster, Graham, Rogier Creemers, Paul Triolo, and Elsa Kania. "Full Translation: China's 'New Generation Artificial Intelligence Development Plan' (2017)." DigiChina project (blog). New America Cybersecurity Initiative, August 1, 2017. <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.

- Xinxin, Zhang, and Chen Weiwei. “Commissioner Li Yanhong: The Era of Artificial Intelligence Should Break ‘Data Islands’ and ‘Innovation Islands.’” [In Chinese.] Xinhua News Agency, March 2, 2019. <http://lianghui.people.com.cn/2019cppcc/n1/2019/0302/c425500-30953752.html>.
- Zwetsloot, Remco, and Dahlia Peterson. “The US-China Tech Wars: China’s Immigration Disadvantage.” *Diplomat*, December 31, 2019. <https://thediplomat.com/2019/12/the-us-china-tech-wars-chinas-immigration-disadvantage/>.
- Zwetsloot, Remco, Helen Toner, and Jeffrey Ding. “Beyond the AI Arms Race.” *Foreign Affairs*, November 16, 2018. <https://www.foreignaffairs.com/reviews/review-essay/2018-11-16/beyond-ai-arms-race>.
- Zwetsloot, Remco, Jacob Feldgoise, and James Dunham. *Trends in U.S. Intention-to-Stay Rates of International Ph.D. Graduates across Nationality and STEM Fields*. Washington, DC: Center for Security and Emerging Technology, Georgetown University, April 2020. <https://cset.georgetown.edu/wp-content/uploads/CSET-Trends-in-U.S.-Intention-to-Stay-Rates.pdf>.

About the Author

Christine Fox is currently the assistant director for policy and analysis at the Johns Hopkins University Applied Physics Laboratory, where she is responsible for connecting APL's considerable technology expertise to broader policy issues. Previously, she served as acting deputy secretary of defense between December 2013 and May 2014, and from November 2009 until July 2013, Ms. Fox served as the director, cost assessment and program evaluation in the Office of the Secretary of Defense. She also formerly served as the president of the Center for Naval Analyses, a federally funded research and development center. In addition to her position at APL, she is on the Board of Directors of the US Naval Institute, a trustee for the Woods Hole Oceanographic Institution, a member of the Advisory Committee for the National Academy of Sciences Division on Engineering and Physical Sciences, and is a life member of the Council on Foreign Relations.



JOHNS HOPKINS
APPLIED PHYSICS LABORATORY