



CAN DEMOCRACIES COOPERATE WITH CHINA ON AI RESEARCH? REBALANCING AI RESEARCH NETWORKS

CAMERON F. KERRY | JOSHUA P. MELTZER | MATT SHEEHAN

每日免费获取报告

1、每日微信群内分享7+最新重磅报告;

2、每日分享当日华尔街日报、金融时报;

3、每周分享<mark>经济学人</mark>

4、行研报告均为公开版,权利归原作者所有,起点财经仅分发做内部学习。

扫一扫二维码 关注公号 回复:研究报告 加入"起点财经"微信群。。



Authors

Cameron F. Kerry is the Ann R. and Andrew H. Tisch distinguished visiting fellow at The Brookings Institution, where he is a frequent speaker, writer, and commentator on public issues involving information technology. In addition to his work at Brookings, Kerry is also a visiting scholar with the MIT Media Lab. Previously, he served as General Counsel and Acting Secretary of the U.S. Department of Commerce, where he led the Obama administration's work on consumer privacy, including the Consumer Privacy Bill of Rights. He has also spent time in private practice as a litigator and regulatory lawyer.

Joshua P. Meltzer is a senior fellow at The Brookings Institution, where his research focuses on digital trade, emerging technologies, and Al. Meltzer has testified before the U.S. Congress, the U.S. International Trade Commission, and the European Parliament. He has been an expert witness in litigation on data flows and privacy issues in the EU and a consultant to the World Bank on trade and privacy matters. Meltzer is a member of Australia's National Data Advisory Council. He has taught digital trade law at Melbourne University Law School and at the University of Toronto Law School, where he was an adjunct professor, and taught ecommerce and digital trade at the UK Foreign, Commonwealth and Development Office diplomatic academy. Prior to joining Brookings, Meltzer was posted as a diplomat at the Australian Embassy in Washington D.C. and before that was a trade negotiator in Australia's Department of Foreign Affairs and Trade.

Matt Sheehan is a fellow at the Carnegie Endowment for International Peace, where his research covers global technology issues with a focus on China. His research projects explore China's artificial intelligence ecosystem, the future of Chinese technology policy, and the role of technology in China's political economy. Matt is the author of *The Transpacific Experiment: How China and California Collaborate and Compete for our Future* (Counterpoint Press, 2019). From 2010-2016 Matt lived and worked in China, including as the first China correspondent for the World Post. Matt worked as a fellow at the Paulson Institute's think tank, MacroPolo, where he led research on Chinese technology issues.

Acknowledgements

The authors are grateful for the research assistance of Mishaela Robison and Aaron Tielemans at the Brookings Institution. Aaron Tielemans also acted as project manager for FCAI throughout the evolution of this paper, and Andrea Renda of the Centre for European Policy Studies. The authors are also thankful for the editing and production assistance from Izzy Taylor, Antonio Saadipour, and Esther Lee Rosen of the Brookings Institution and peer review from Ryan Hass, Henry Gao, and Kevin Wolf.

The Brookings Institution is a nonprofit organization devoted to independent research and policy solutions. Its mission is to conduct high-quality, independent research and, based on that research, to provide innovative, practical recommendations for policymakers and the public. The conclusions and recommendations of any Brookings publication are solely those of its author(s), and do not reflect the views of the Institution, its management, or its other scholars.

Amazon, Apple, Facebook, Google, IBM, and Microsoft are donors to the Brookings Institution.

Brookings is committed to quality, independence, and impact in all of its work. Activities supported by its donors reflect this commitment.

Introduction

China looms large in the global landscape of artificial intelligence (AI) research, development, and policymaking. Its talent, growing technological skill and innovation, and national investment in science and technology have made it a leader in AI.

Over more than two decades, China has become deeply enmeshed in the international network of AI research and development (R&D): co-authoring papers with peers abroad, hosting American corporate AI labs, and helping expand the frontiers of global AI research. During most of that period, these links and their implications went largely unexamined in the policy world. Instead, the nature of these connections was dictated by the researchers, universities, and corporations who were forging them.

But in the past five years, these ties between China and global networks for R&D have come under increasing scrutiny by governments as well as universities, companies, and civil society. Four factors worked together to drive this reassessment: (1) the growing capabilities of AI itself and its impacts on both economic competitiveness and national security; (2) China's unethical use of AI, including its deployment of AI tools for mass surveillance of its citizens, most notably the Uyghur ethnic group in Xinjiang but increasingly more widespread; (3) the rise in Chinese capabilities and ambitions in AI, making it a genuine competitor with the U.S. in the field; and (4) the policies by which the Chinese state bolstered those capabilities, including state directed investments and illicit knowledge transfers from abroad.

Taken together, these concerns led to intense scrutiny and new questions about these long-standing ties. Is cooperation helping China overtake democratic nations in AI? To what extent are technologists and companies in democratic nations contributing to China's deployment of repressive AI tools?

This working paper considers whether and to what extent international collaboration with China on AI can endure. China has been a subject of discussions among the government officials and experts participating in the Forum for Cooperation on AI (FCAI) over the past two years. The 2021 FCAI progress report identified the implications of China's development and use of AI for international cooperation.¹ The report touched on China in connection with several of the recommendations regarding regulatory alignment, standards development, trade agreements, and R&D projects but also focused on Chinese policies and applications of AI that present a range

of challenges in the context of that nation's broader geopolitical, economic, and authoritarian policies. A roundtable discussion on December 8, 2021 presented these issues to FCAI participants more fully and elicited their views.

This paper expands and distills this work with a focus on the scope, benefits, and prospective limits of China's involvement in international AI R&D networks. In Part I, it presents the history of China's AI development and extraordinarily successful engagement with international R&D and explains how this history has helped China become a global leader in the field. Part II shows how China has become embedded in international AI R&D networks, with China and the United States becoming each other's largest collaborator and China also a major collaborator with each of the other six countries participating in FCAI. This collaboration takes place through multiple pathways: enrollment at universities, conferences, joint publications, and work in research labs that all operate in various ways to develop, disseminate, and deploy AI.

Part III then provides an overview of the economic, ethical, and strategic issues that call into question whether such levels of collaboration on AI can continue, as well as the challenges and disadvantages of disconnecting the channels of collaboration. The analysis then looks at how engagement with China on AI R&D might evolve. It does so primarily through a U.S.-focused lens because the U.S., as by far China's largest competitor and collaborator in AI, provides an umbrella and a template for countries and FCAI participants that also collaborate with China on AI R&D and face many of the same issues. Moreover, measures to respond to the challenges China presents are more likely to be effective in coordination than in isolation. Recent U.S. export controls on semiconductors and the technologies used to manufacture them have laid bare the critical role of countries such as Japan and Korea. For now, the U.S. government is able to force foreign compliance through administrative measures, such as the foreign direct product rule, but these mechanisms may be made moot if foreign manufacturers engineer U.S. technology out of their supply chain. This paper deals with cooperative research rather than hardware supply chains, but similar dynamics exist across these domains. Accordingly, this paper is not just about collaboration with China but also about collaboration in relation to China.

The U.S., other governments participating in FCAI, and their partners are not the only actors in this drama. What AI R&D with China looks like going forward will also be determined by what China does. China's intensifying push for technological self-reliance has accelerated China's disengagement from the international technology ecosystem in certain respects, while so far keeping it deeply enmeshed in other international research networks. The future

Measures to respond to the challenges China presents are more likely to be effective in coordination than in isolation. trajectory of this engagement will depend heavily on actions taken by the Chinese government and the Chinese Communist Party.

In light of the issues presented by these changes, the paper proposes rebalancing AI R&D with Chinese researchers and institutions through a riskbased approach. Going forward, such collaboration will require a clear assessment of the costs and benefits, aiming to maximize the benefits of an open research environment and strong international links with the risks presented by AI R&D with China. Adopting an appropriately risk-based approach often will not counsel complete disengagement with China on AI R&D and instead require a rebalancing that takes into account the various vectors for knowledge transfer. Crucially, governments need to work collaboratively with each other and with companies, universities, and research labs to inform the assessment of the risks and understand the benefits of AI R&D with China. A failure to build these partnerships into the risk-assessment process could lead to bad outcomes that mismeasure risks and benefits, leaving the U.S. worse off.

The table below summarizes for each pathway of collaboration discussed in Part II (1) the costs and benefits discussed in Part III and (2) the recommendations for risk mitigation strategies and international cooperation needed for the rebalancing in each network to be globally effective. The benefits, risks, and challenges vary for each pathway of collaboration; as a result, enrollment of foreign students and dissemination of knowledge through publication and conferences should continue with limited exceptions, while active research collaborations call for more significant restraints on the part of individual researchers and institutions, informed and guided by governments.

CAN DEMOCRACIES COOPERATE WITH CHINA ON AI RESEARCH? | REBALANCING AI RESEARCH NETWORKS

AI R&D Networks	Costs	Benefits	Risk Mitigation Strategy	Role for International
Joint U.SChina Al publications	Exposure to knowhow (beyond what is gained from merely reading the publication) that can be used to advance China's AI capability in ways that is detrimental to U.S. values, economics, or national security	Access to Al knowledge in China that supports U.S. innovation Opportunity to understand the pace and focus of Al R&D in China Research relations can lead to graduate study at U.S. universities	Develop an approach to screening joint research that takes into account whether it provides Chinese researchers with new information that they would not obtain on their own; whether the Chinese partner is affiliated with the Chinese military (i.e., Seven Sons of National Defense); and taking account of the potential applications of the research	Alignment with key Al powers (EU, Japan, U.K., and others) is needed to prevent U.S. action merely leading to Chinese researchers increasing joint publications with Al researchers in third countries
International AI conferences	Chinese researcher access to cutting-edge AI developments in the U.S.	Provides U.S. researchers with insights into AI developments in China Opportunity to understand pace and focus of AI R&D in China Networking	Focus narrowly on screening out bad actors from conferences to limit access to the types of AI knowledge and networks gained from participating in conferences.	Top AI conferences are global, requiring international alignment
Enrollment in computer science and other Al- adjacent programs at U.S. universities	Hands-on training of technical researchers and engineers, many of whom will stay in the U.S. but some of whom may return to China	Economic gains for the university Opportunities for research that benefit students/faculty Exposing Chinese AI talent to U.S. values and opportunity	Principles on joint Al research to guide universities on how to screen out high risk students (e.g., from the CCP) Work with law enforcement to identify bad actors Increase opportunity for Chinese students/researchers to	Alignment with other universities in the West needed to address leakage (i.e., high-risk Chinese students studying at non- U.S. universities) International cooperation on a common approach to research principles

			remain in the U.S. post- graduation (i.e., provide a clear path to permanent residency and citizenship)	
Open publication of AI research	Access to cutting-edge Al research	Access of U.S. researchers to Al developments in China	Curtailing open publication likely not feasible and inconsistent with broader U.S. support for an open internet	
U.S. Private sector labs in China/China private sector labs in the U.S.	Training and access to commercially sensitive AI knowledge Requirements for technology transfer	Close engagement with top Chinese AI researchers Increased understanding of Chinese AI ecosystem Test and trial AI products in China to measure commercial impacts	Limits on types of Al research that can be conducted Cooperation with U.S. law enforcement to address risk of theft of Al knowledge Support international rules on technology transfer requirements	Work with partners to agree to rules prohibiting technology transfer requirements (e.g., non source code requirements in trade agreements) International cooperation to address risk of IP theft

We also recommend establishing an institution that can facilitate information sharing amongst government, academia, and civil society on risks from international AI collaboration with China. This information would inform the risk-based analysis that will need to underpin any rebalancing of AI R&D networks with China.

China's role in international AI R&D

"Science is the most international field in the world today," the editor of Scientific American said last year. International scientific collaboration affords opportunities to share project costs, resources, and expertise. These opportunities are especially important when it comes to AI R&D because, for reasons described in the FCAI progress report, R&D in AI is especially collaborative. Moreover, joint research can have higher impact, as the outputs are more accepted globally, and it provides greater access to new ideas.² Indeed, research shows that scientific publications that are the result of international collaborations are more highly cited than less collaborative ones.³ As a result of all these factors, collaboration in AI R&D across disciplines, organizations, and international borders is extensive.

China has made the most of the system of international scientific collaboration through collaborative research, study abroad, monitoring of international publications, hosting foreign technology labs, and technology transfer by a variety of means.⁴ This section explores how this has especially been the case with respect to R&D on AI, helping China to become a global leader that is deeply embedded in the R&D ecosystems of the U.S. and other AI leaders around the world.

25 years of Chinese AI development

A full exploration of how China has become a leader in AI requires a brief review of the history of Chinese AI development and what role international cooperation has played in that process. The following outlines how Chinese researchers and AI developers have benefited from all the ways that AI R&D occurs, with a focus on studying abroad, online publication of AI papers, joint publication with overseas researchers on AI, and employment in U.S.-owned AI labs in China.

The Chinese Communist Party has long taken a techno-nationalist view of state power.⁵ But in the nearly three decades of Mao Zedong's reign, the country's adversarial relationships with the world's leading scientific powers forced China to seek technological strength from within, leading to a highly isolated science and technology ecosystem. But with the beginning of China's "reform and opening" period in the late 1970's, China began to systematically engage with the international science and technology community as a means of absorbing the key technologies and knowhow needed to modernize China's

economy and military. That international engagement included sending Chinese scholars abroad, making foreign technology transfer a prerequisite for market access, systematically monitoring global scientific progress, and outright theft of intellectual property, among other tactics.

China's pursuit of AI development has followed some of these same trajectories, but it has also been driven heavily by more organic academic and business exchanges. Researchers in China have been working on AI for decades, but for most of that time, they remained far behind the global cutting edge. China's Cultural Revolution (1966-76) dramatically weakened the country's intellectual and scientific community, and its knock-on effects continued strongly through the 1990s. Chinese universities lacked the infrastructure for cutting-edge research, and China had few technology companies with the resources to pour into knowledge-intensive R&D.

The country's AI ecosystem got a boost in 1998 when Microsoft CEO Bill Gates decided to open a research lab in Beijing. The lab, known as Microsoft Research Asia (MSRA), was led by Kai-Fu Lee, a Taiwan-born AI researcher who had worked on cutting-edge speech recognition technology at Apple in California before joining Microsoft. MSRA quickly became a hotbed for AI work and a magnet for talented Chinese researchers. Many of those researchers would go on to lead technical teams at China's internet giants and create some of China's most influential AI startups.⁶ By 2004, the MIT Tech Review had crowned MSRA "the world's hottest computer lab."⁷

Despite that early promise, Chinese progress in AI and related fields still flew under the radar of most foreign governments. The real-world use cases of AI remained limited, and Chinese technology companies were not viewed as true competitors with their international peers. Over the following decade, both of those things changed.

Beginning around 2012, several factors converged to accelerate China's Al development. Deep learning—a machine learning (ML) technique for constructing more powerful neural networks—exploded, opening up countless new applications in everything from autonomous vehicles to facial recognition and other imaging. As one of the few countries to foster its own homegrown internet giants—home to large pools of engineering talent, training data, and computing power—China was in a strong position to advance deep learning and apply it to existing products as well as further research.⁸ China's vast population provided sources of data on the massive scale needed to train these neural networks, and China's internet companies were able to exploit this data to develop and deploy new products on a scale that enabled them to become powerful and innovative national champions.

China's vast population provided sources of data on the massive scale needed to train these neural networks, and China's internet companies were able to exploit this data to develop and deploy new products on a scale that enabled them to become powerful and innovative national champions. Chinese private sector interest in AI continued to grow through the mid-2010s, with high-quality research ramping up and dozens of new startups applying the technology to different market niches. Much of that excitement was homegrown, but it also frequently intersected with, or was inspired by, the remarkable achievements of international AI labs. Some of China's top AI startups in computer vision, autonomous vehicles, and speech recognition were founded by technologists who had returned home after studying or working abroad.⁹ Overseas, Baidu turned heads in 2014 when it hired Andrew Ng,¹⁰ a deep learning pioneer at Stanford and founder of the Google Brain team, to be its chief scientist. Ng helped build out Baidu's footprint in Silicon Valley and bolster its reputation as a hub for AI research.

With momentum building in Chinese companies and AI becoming a mainstream technology worldwide, the Chinese government formally entered the fray in July 2017. The State Council released its New Generation Artificial Intelligence Plan,¹¹ which laid out high-level goals for the country's AI capabilities, including making China "the world's primary AI innovation center" by 2030. This ambitious target led many in the U.S. and European business communities to begin taking China seriously as a competitor in the field.

The most direct impacts were felt within China, where the plan acted as a tremendous catalyst for AI activity by all aspects of the Chinese bureaucracy and business community.¹² Local officials subsidized AI startups, purchased AI-powered products, and partnered with technology companies on pilot projects. Private investment surged in China's AI industry,¹³ and firms across the country rebranded themselves as "AI companies." Much of this spending turned out to be wasted and many of these projects failed, but the sheer amount of new activity also proved effective in rapidly growing China's AI ecosystem and giving its scientists and companies the resources and runway to experiment with new models.

Looking to piggyback on the momentum in China, several leading Al institutions followed in Microsoft Research's footsteps by straddling the U.S. and Chinese AI ecosystems, drawing on technical talent and funding in both countries. In December 2017, Google announced the creation of a new AI lab in Beijing,¹⁴ and in 2018, the Massachusetts Institute of Technology formed a five-year collaboration with Chinese speech recognition giant iFlyTek.¹⁵ At least in the case of Microsoft, this strategy continued to yield major dividends in terms of research output. According to one 2018 study, the Microsoft Research Asia lab produced nearly one third of Microsoft's most-cited AI research papers globally.¹⁶

Comparing and contextualizing Chinese AI capabilities

Where does China stand today in terms of AI capabilities? A full accounting of comparative capabilities is beyond the scope of this brief, but a survey of key indicators provides a sketch. We begin by dividing AI capabilities into two categories: commercial applications and basic research, with the latter having more straightforward methods of measurement.

Looking at commercial applications, in 2021, China ranked second behind the United States in both newly funded AI companies and private sector AI investment.¹⁷ Zooming in on top AI startups, as of 2020, China accounted for three of the top 10 most valuable AI startups globally,¹⁸ with the U.S. accounting for the other seven. Two of those Chinese startups—Megvii and SenseTime—are computer vision companies that have been blacklisted by the U.S. Commerce Department for their alleged role in human rights violations in Xinjiang.¹⁹ Looking at AI patents,²⁰ of the top 20 companies by AI patents granted, 10 are headquartered in the U.S. and five are headquartered in China, with the remaining five in Japan, South Korea, Germany, and the Netherlands.

Though all of these indicators give some sense of national capabilities in commercial AI applications, they are only small windows into a far more complex reality: tens of thousands of companies directly or indirectly apply AI to upgrading their businesses. A best guess description of that reality might be to say that the United States likely leads the world in most commercial AI applications but that China is not far behind on the whole and leads in certain categories. Many of China's most common AI applications are currently in surveillance and public security, but some of these are beginning to shift toward broader uses of the technology to upgrade China's industrial and agricultural economies.^{21,22}

Looking at AI research capabilities, China has risen to match or exceed the U.S. in the *quantity* of *high-quality* research but still lags behind the U.S., and arguably the EU and U.K., in the kinds of paradigm-shifting breakthroughs that have reshaped the field. The most common way to assess research capabilities is the publication of academic papers, often using either citation totals or acceptance at elite AI conferences as a filter for quality. Looking at citations, one study by the Center for Security and Emerging Technology (CSET) looked at the top 1% most-cited AI research papers,²³ finding that China surpassed the EU in 2016 and the U.S. in 2019 (data for 2020 and 2021 was not published). Looking at different subfields of AI, that study found China leading substantially in computer vision research, trailing the U.S. by a wide margin in natural language processing and essentially tying with both the U.S. and EU in robotics research. These numbers are helpful for assessing the broad base of high-quality AI research output, but the huge quantity of papers

in the sample (1.73 million AI papers over 10 years) means that even the top 1% still totaled over 17,000 academic papers, many of which are highly cited but not necessarily "game-changing" in any given field.

Another way to measure high-quality AI research output is to look at which papers gain acceptance to the most elite and selective AI conferences. Studies using AI conference acceptances may give a better picture of the truly elite research papers in a given year, but they also introduce new potential biases to the selections. In different studies of top AI conferences,^{24,25} the United States and the European Union perform better, with the U.S. usually leading by a wide margin, and China often on par with the EU. These studies classify researchers by their current affiliation, but it is worth noting that researchers originally from China comprise a very large portion–27% according to one study–of those doing this elite AI research at American institutions.²⁶

When it comes to game-changing breakthroughs—the rare advances that substantially alter the direction or trajectory of the field—most observers both in and outside of China agree that the majority of these come from the U.S. and a handful of democratic allies: Canada, the U.K., and the EU. The pioneers of deep learning all live and work in Canada and the U.S., the same countries where key innovative follow-ups such as recurrent neural networks and generative adversarial networks were developed. More recently, major advances such as transformers and novel applications like large language models have also been pioneered in the United States. Researchers in China have contributed some game-changing breakthroughs of their own—one of which will be described later in this paper—but thus far have tended to be fast followers when it comes to these types of advances.

When it comes to game-changing breakthroughs—the rare advances that substantially alter the direction or trajectory of the field—most observers both in and outside of China agree that the majority of these come from the U.S. and a handful of democratic allies: Canada, the U.K., and the EU.

None of these single studies can definitively describe the state of China's AI research capabilities, but, taken together, they resonate with the broad assessments of many leading AI researchers in both China and the United

States: while China produces a huge amount of very good AI research, American companies and universities are still the driving force behind most of the game-changing breakthroughs of recent years.

Tensions grow

The development of China into a global leader on AI will affect AI development globally, including how AI is used and in what context and the values that AI systems reflect. This impact of China's AI capacity is also happening at a time of growing strategic competition with the U.S., EU, and other democratic countries that are leaders in AI. The dual-use nature of AI also means that leadership in AI will strengthen military capacity and raise national security concerns.

These developments, paired with the PRC government's demonstrated willingness to use AI in unethical ways, are of particular concern for democratic nations. This has been particularly the case with China's use of AI to boost its surveillance capacity. The Chinese government had long made surveillance of its population a high priority, but for decades, the actual identification of individuals required human eyes and ears to understand the materials. The advent of more accurate facial and speech recognition changed that requirement, allowing computer vision algorithms to sift through thousands of hours of footage to identify or track individuals. This phenomenon took hold most dramatically in China's western region of Xinjiang, where the government was executing a sweeping campaign of mass detention and cultural repression targeting Uyghurs and other Muslim minorities. Beginning around 2016-2017, Xinjiang turned into a "frontline laboratory for surveillance."²⁷ The public security apparatus began using AI to identify and monitor Uyghurs,²⁸ over a million of whom are estimated to have been detained in "re-education camps."²⁹ Elsewhere in China, facial recognition was used to catch alleged criminals at beer festivals and pop concerts and to build "smart cities" that optimized traffic flows.^{30,31} Much of the deployed surveillance technology turned out to be less high-tech than advertised-often lacking an actual AI backend-but its omnipresent nature and accompanying repression gave a glimpse into the dangerous possibilities of an authoritarian state with access to AI tools.

As Chinese AI capabilities have grown, U.S.-China tensions have deepened,³² ethical questions mounted,³³ and international AI R&D began to fray. For instance, after less than two years, the MIT-iFlyTek partnership was canceled over human rights concerns.³⁴ By 2020, the big ambitions for the Google AI lab in Beijing had largely fizzled out, and in 2021, IBM closed a long-running

research lab in Beijing.³⁵ And in spring of 2022, Microsoft Research Asia reportedly stopped accepting interns from Chinese universities with deep ties to the military.³⁶ At a macro scale, co-authorship of AI research papers between U.S. and Chinese institutions peaked in 2019,³⁷ though it remains unclear whether this is a short-term disruption or indicative of an enduring change.

Chinese students and researchers in the U.S. have also increasingly come under government scrutiny. In 2018, President Trump and his advisers reportedly debated whether to ban all Chinese students from entering the U.S. but eventually decided against the proposal.³⁸ Instead, the Trump administration shortened the duration of visas for Chinese graduate students in certain STEM fields,³⁹ causing many Chinese students to be stranded back in China for months or even years while waiting on those renewals. That same year, the U.S. Department of Justice launched the "China Initiative," a program intended to crack down on economic espionage and other forms of technology transfer, particularly in universities. Though it notched a few major convictions, the program also resulted in the wrongful indictment and eventual acquittal of prominent Chinese and Chinese American professors.^{40,41} A 2021 study found that 42% of scientists of Chinese descent in the U.S. reported feeling racially profiled by the U.S. government,⁴² with many of them choosing to curtail projects that involved Chinese collaborators. In the spring of 2022, the Department of Justice terminated the China Initiative.43

Meanwhile, in China, the government began ratcheting up its control over Chinese technology companies, including some of its AI juggernauts. In April 2021, the Chinese government took a 1% stake and a board seat in the key Chinese corporate entity for Bytedance,⁴⁴ owner of China's leading AI-driven social media app Douyin. These increases in government influence over management of private, AI-focused companies, coupled with a series of regulations on the development and deployment of AI algorithms, portend an emerging AI governance framework in China that has altered the political and ethical calculus around potential international AI collaboration with Chinese entities.

AI R&D networks and pathways of collaboration

The following outlines global networks of AI R&D collaboration and China's role within it. As China's trajectory of development in AI shows, China utilized multiple AI R&D networks to acquire and develop AI knowhow. These pathways of collaboration work independently and together to develop, disseminate, and deploy knowledge about AI. For instance, the huge growth in Chinese students studying postgraduate computer science courses at U.S. universities provided knowledge as well as opportunities to build relationships with U.S. researchers that in turn led to joint publications and employment in academia and the private sector. U.S. companies that opened AI labs in China hired Chinese talent, including some with degrees from U.S. universities, who went on to work at Chinese companies to support their own AI R&D efforts.

The often blurry and sometimes shifting boundaries between the PRC government and Chinese universities, research labs, and companies further complicate potential AI R&D with partners in China should continue.

As outlined, the economic, security, and moral tensions with China over its development and use of AI require assessing how AI R&D networks work and to what extent AI R&D collaboration with China should continue. The often blurry and sometimes shifting boundaries between the PRC government and Chinese universities, research labs, and companies further complicate potential AI R&D with partners in China. And Russia's invasion of Ukraine, paired with the earlier Russia-China joint statement on cooperation in areas including AI, has further raised the stakes on these questions.

Any rebalancing of AI R&D with China needs to consider all of these interconnected AI R&D networks and the interactions among them holistically and globally. Moreover, the sprawling and diverse networks of AI R&D offer numerous paths for disseminating knowhow across the globe. As a result, should the U.S. or any other country act unilaterally to rebalance AI R&D with China, leakage will occur in the absence of international cooperation and alignment. This section looks more closely at the various pathways of collaboration and the extensive role of China and Chinese researchers in these networks.

Enrollment at universities

International students are now significant majorities across all graduate level science and engineering programs at U.S. universities. In the computer sciences, international students accounted for 72% of full-time students in 2019, up from 62% in 2009. Figure 1 shows that international students

comprise a large majority of graduate students in computer and information sciences at U.S. institutions.



Figure 1: Computer and information sciences: Full-time U.S. and international graduate students, 1998-2019

Source: National Science Foundation, survey of graduate students and postdoctorates, National Foundation for American Policy calculations. Note: U.S. students include lawful permanent residents.

These international students contribute the ideas, knowhow, and workforce that make the U.S. a center of innovation and scientific research.⁴⁵ The presence of international student enrollment supports interconnection between overseas and U.S. research communities, including in AI. Industry relies heavily on these international students as a source of high-level technical talent.

The 2019 National Science Foundation's Science and Engineering Indicators for international graduate student enrollment in 2017 and 2018 respectively showed 14,680 and 16,990 computer science students from China.⁴⁶ In both, only India sent more computer science graduate students to the U.S., and only Chinese engineering graduate students outnumbered Chinese computer science students. At the doctoral level, China leads in the cumulative number of U.S. doctorate recipients in computer sciences from 2001 to 2020 with 6,408.

This diaspora of technical talent from China made a major impact on top-tier AI research in the U.S. and around the world. In one study looking at researchers invited to present their research at an elite global AI conference, nearly one third of all the researchers at the conference had completed their undergraduate studies in China.⁴⁷ This showed China to be the largest global source of top-end AI research talent, followed by the United States, the European Union, and India. This talent has had an enduring impact on AI R&D outside of China: the majority (54%) of these Chinese researchers chose to attend graduate school in the U.S., and the country proved to be a sticky location for them. The large majority (88%) of these Chinese researchers who attended graduate school in the U.S. went on to work in the U.S. and publish cutting edge research there, as shown in Figure 2.

Figure 2: More of China's top AI talent ends up in the U.S. than anywhere else



Of 128 researchers with undergraduate degrees from Chinese universities whose papers were presented at the A.I. conference, more than half now work in the U.S.

By Ella Koeze • Original sample was made up of 671 authors of a random selection of 175 papers selected from the over 1,400 papers presented at NeurIPS 2019, a top A.I. conference. None of the 128 researchers represented here are current students. Post-graduate work countries are based on where the researcher lives, not where their company or institution is headquartered. Data is current as of the first quarter of 2020.

Source: New York Times from data in The Global AI Talent Tracker (MacroPolo, 2020).⁴⁸

This pipeline of research talent may attenuate, however. More recent data on enrollment of Chinese students in graduate education shows significant declines. Again, according to NSF data, Chinese student enrollment in computer science in the U.S. was at 14,780 in 2020,⁴⁹ a 15% decline. This decline likely largely reflects the impact of COVID-19 and restrictions on travel.

However, further declines are expected due to ongoing COVID-19 restrictions in China, as well as the impact from the Trump-era Presidential Proclamation (PP10043), which restricts access to U.S. universities of Chinese graduate students who receive funding from or who have been employed by, studied at, or conducted research on behalf of an entity that implements the China military-civil fusion strategy, a potential very broader range and number of students.⁵⁰

A likely secondary effect of PP10043 will be to chill application to the U.S., increasing application by Chinese graduate students to universities in the U.K., Canada, Australia, to name some alternative research leaders.⁵¹ Other countries' own concerns about China, especially when it comes to AI, and China's policies may diminish this flow of Chinese students more broadly. Given the scale of Chinese international study, the flow will almost certainly continue at significant levels even in computer science and study related to AI.

Co-authorship of AI publications

Publication of research results is a central feature of higher education, especially at the graduate level. The involvement of international students in the pipeline of researchers, as well as wide interest in AI across many countries and research institutions, serves to broaden AI R&D across authors of multiple nationalities and affiliations. These collaborations operate to disseminate knowledge both through participation and through publication. Such collaborations can elevate the quality of the research itself by involving the best available talent to sharpen and improve the research product. Empirical studies have found that papers co-authored by researchers from different countries receive significantly higher citation counts.⁵²

In terms of total output of AI research papers in 2020—not filtered for quality of publications—China leads all other countries and the EU (with the U.K. still included in the latter) (Figure 3). While collaboration makes up a substantial share of that output (12%), China would remain the leader in total publications even if all of its international collaborations were removed from this data. Second and third in total research output are the EU and U.S., both of which collaborate internationally on approximately half of all papers. Compared with other FCAI countries, Japan appears to contain a relatively insular AI research community, while Australia, Canada, and Singapore all see high levels of international collaboration with both China and FCAI countries.





Source: Emerging Technology Observatory's County Activity Tracker: Artificial Intelligence⁵³

When it comes to collaboration between China and other AI leaders, U.S.-China collaborations led the way, but AI research collaboration with China was by no means a uniquely American phenomenon. Appendix A shows the levels of collaboration between each country (or union) participating in FCAI with each other as well as with China. The tables in Appendix A show that the U.S. is the top collaborator for Australia, Canada, China, the EU, and Japan. For the U.S. and Singapore, China is the top partner, with Chinese co-authors contributing to just over one third of all of Singapore's AI research papers and appearing on 15% of all U.S. AI publications.



Figure 4: Country co-authorship over time (United States)





Figure 5: Country co-authorship over time (China)

Source: Emerging Technology Observatory's Country Activity Tracker: Artificial Intelligence⁵⁵

Between 2010 and 2021, the number of AI research papers that include an author in both China and America more than quadrupled, rising to 9,458 in 2021.⁵⁶ This is triple the number of research papers from the next most frequent country-to-country collaborators, the United Kingdom and the United States. The massive scale of U.S.-China collaborative research reflects that

the two countries are the largest producers of AI research, as well as the connections forged by individuals who have worked or studied in both.

Chinese researchers also collaborated with their peers across the globe, ranking in the top three partners for many countries alongside the United States and the EU. The additional tables in Appendix show the total volume of AI research papers from each of the FCAI countries (the U.S., Australia, Japan, Singapore, Canada, and the EU) from 2010 to 2021 and the shares of those papers that involve international collaboration with each other FCAI participant.

Taken together, the data reveals AI R&D networks that are highly international and deeply interconnected among all AI leaders.

Online publication of AI research

The explosion of progress in the AI field during the past 10 years and global dissemination of AI knowledge (regardless of whether it was developed collaboratively or not) has been accelerated by an evolution of research norms toward early publication online in open-source repositories. Instead of waiting through the extended peer review to print new research in expensive academic journals or present them at exclusive conferences, AI researchers often immediately upload "pre-prints" of their academic articles to open online repositories like arXiv (pronounced "archive") that do not require peer review.⁵⁷ Free and publicly available pre-prints of AI papers have gained popularity, with over 50,000 such papers uploaded in 2021 alone.⁵⁸

This rapid open-source dissemination is particularly important in AI subfields such as computer vision and natural language processing, where new performance records are often set on a monthly or even weekly basis, incentivizing scientists to document their moment at the global cutting edge. It allows them to immediately plant a stake for results that break existing records for AI performance on established benchmarks or novel ways of processing training data or wiring neural networks.⁵⁹ In many cases, the authors of these papers even include the code they used to implement the ideas from the paper,⁶⁰ as well as the data it was trained on, making replication even easier.

Even for more deliberate and peer-reviewed academic AI conferences, all papers selected eventually are posted publicly on the conference websites. For example, any AI researcher around the globe can gain access online to over 30 years of papers from one of the top conferences on neural networks and machine learning.⁶¹ In recent years, several major ML conferences have

begun asking authors to either include the code they used or justify why they aren't including it.⁶²

There are some areas of research involving process knowledge where more hands-on training is required. These often focus on hardware-heavy aspects of research, such as the construction of computer clusters for training large AI models. But many AI research papers published today contain the full scope of new knowledge generated and are rapidly digested and applied by researchers and engineers around the world. The result is a global accelerated feedback loop for AI research, one in which researchers across the globe are both contributing to and building off of each other's work in real time.

There are some notable exceptions to the norm of open publication. Within companies, when R&D moves from the research stage into actual product development, the work often remains proprietary. Similarly, when AI scientists do highly applied research for military or government clients, that work is subject to existing and often highly restrictive limitations on both who can do the work and how it can be described in public.

There is also an emerging—and highly controversial—trend in some circles away from open-source release of AI R&D that is worth noting here. This is a limited move toward "staged release" or "limited release" of certain very large or powerful AI models to limit dissemination of otherwise open foundational research.

The most prominent example is the choice by the leading research lab OpenAI in February 2019 not to release the full version of its most novel and powerful neural network language model, GPT-2. Instead, it released a much smaller version of the model.⁶³ Over the next nine months, OpenAI steadily released larger versions of the model while researchers around the world explored ways the model might be misused by bad actors (such as generating fake news or automating abusive content on social media) and ways of mitigating these harms. Once it was satisfied, they could be mitigated, and OpenAI finally released the full 1.5 billion parameter model. Yet, when OpenAI released its follow-up model, GPT-3, it chose to simply describe the model in a research paper and then offer limited access for paying users.⁶⁴

Staggering the release of GPT-2 and not publicly releasing GPT-3 were controversial decisions within the AI community. Many leading researchers criticized the move as a publicity stunt, highlighting just how strong the push for open publication and access remains.⁶⁵ In response, Meta's AI research team decided to build a comparable large language model—called OPT—and in May 2022 the company released it for free to the public.⁶⁶ Some other researchers, however, saw value in experimentation with forms of staged or

limited release of large models that have the potential for misuse by bad actors. In a major paper on very large and powerful AI models (often called "foundation models"),⁶⁷ dozens of researchers at the Stanford Institute for Human-Centered Artificial Intelligence outlined recommendations on how staged-release programs could be tied to independent auditing of models to predict and mitigate misuse of the technology.

Conferences

Conferences provide another pathway for dissemination of information about AI R&D through presentation of research, publication, and exchanges of ideas among participants. The emergence of pre-print publication has reduced some of the significance of presentations, but, as conference attendees in many fields are rediscovering with the re-emergence of in-person events in the wake of some COVID-19 restrictions being lifted, there is added value in personal contacts. They provide social learning, additional information that may not be reflected in publications, and professional networking that may lead to new collaboration or R&D.

Data on conferences provide information about networks of AI R&D: researchers and research teams, subjects of research, and gauge of quality according to selectivity or reputation of the conference. Looking at recent and future locations for 16 major AI conferences (including locations for which conferences were switched to virtual events due to the pandemic),⁶⁸ the United States leads in hosting by a wide margin. Out of 16 AI conferences, the United States is scheduled to, or has hosted, eight at least once and several others multiple times. Of the same grouping, Canada has hosted six (hosting twice for only one conference), and China has hosted five.

Another pair of analyses, examining accepted papers at NeurIPS 2020 and ICML 2020,^{69,70} reflects how conferences can be valuable connectors. These found both the United States and China stood out from other countries in that multiple private companies contributed papers alongside academic institutions. For the U.S., top private institutions include Google, Microsoft, and Facebook, and for China, Tencent, Alibaba, Baidu, and Huawei submitted papers that were accepted to these conferences.

US private sector labs in China/Chinese private sector labs in the US

The foregoing contribution of Chinese and American companies to the NeurIPS and ICML conferences is just one example of the indispensable role that private companies play in AI R&D networks as funders, employers, collaborators, and research hubs. Companies establish AI labs for a variety of reasons, especially access to talent, proximity to customers, access to markets, and cost savings.⁷¹ According to a CSET paper, which focused on the global R&D activities of Amazon, Apple, Facebook, Google, IBM, and Microsoft, the majority of these companies' AI labs (68%) are outside of the U.S., located mostly in Europe (mainly the U.K. and France), as well as Israel and China (approximately 10% of total AI labs in each).⁷² Baidu for the past decade has maintained a significant outpost for the company in Silicon Valley, partly as a way of working with top technical talent based in the United States.

As ethical, competitive, and geopolitical concerns about China and Al development have heightened in recent years, these interconnections have come under increased scrutiny. As ethical, competitive, and geopolitical concerns about China and Al development have heightened in recent years, however, these interconnections have come under increased scrutiny. Some have argued that the national security risks to the U.S. of having U.S. companies with Al labs in China is now too great to justify any presence in China at all and that hostile governments can access technological developments, enabling catch-up with the U.S. on Al.^{73,74} Indeed, as discussed in the section above on China's Al development, Google and other American institutions have closed down initiatives, and several other Chinese technology firms have set up Al labs in the U.S. While several have been quietly shuttered in recent years, some labs in both countries remain.

The existence of such labs creates both thorny tradeoffs in terms of ethics and international technology competition and questions that often don't have clear-cut answers. For a window into what these tradeoffs look like in practice, we examine a case study centered on the most-cited AI research paper of the past decade. We will look briefly at the paper, the institution, and the researchers themselves.

In 2015, a group of researchers at Microsoft Research Asia (MSRA) in Beijing published a paper titled "Deep Residual Learning for Image Recognition"— ResNet for short. ResNet introduced a novel technique for stacking more layers on a neural network, making deep learning "deeper" and turbocharging its accuracy. It accelerated progress in domains such as image recognition, natural language processing, and speech recognition. ResNet's simplicity and wide applicability meant that over the next seven years, it would accumulate over 100,000 citations, making it not just the most-cited paper in AI, but the most-cited paper *in any field* during that time.

Since its founding in 1998, MSRA has been a source of top talent for Microsoft and a training ground for many top Chinese technology entrepreneurs and AI scientists. Of the four scientists who co-authored ResNet, all were educated in China and Hong Kong before joining Microsoft. Since publishing the paper, all four have departed MSRA and taken their talents elsewhere. The celebrated lead author joined Facebook in 2016 and remains there as of 2022. Two others joined Megvii, a computer vision startup which has drawn U.S. sanctions for supporting China's surveillance apparatus in Xinjiang. The fourth co-founded a Chinese autonomous vehicle startup, Momenta, before joining NIO, an electric vehicle firm.

Looking at the people, institutions, and ideas involved in ResNet, we see a mixed bag of beneficiaries and a lot of thorny tradeoffs. While Microsoft may have enjoyed a small first-mover advantage due to in-house development, ResNet's techniques were published openly and rapidly applied by researchers and engineers across the globe. Those applications included everything from major advances in predicting protein structures to the smothering surveillance of ethnic minorities and political dissidents. After publication, the paper's authors went to work for both American and Chinese companies in industries both promising and deeply problematic.

From the perspective of geopolitical technology competition, exploring historical counterfactuals further complicates the picture. What if Microsoft had never opened its lab in Beijing? This would have slowed global AI advances and hurt Microsoft, but it likely would have hurt China far more. Back in 1998, China's local AI ecosystem was practically nonexistent, and without MSRA, it would have taken longer to get off the ground.

But what if concerns over Chinese competition led U.S. policymakers to close MRSA in 2014? The authors might have transferred to Microsoft headquarters, where they would perfect and publish ResNet. They could have left Microsoft and begun working at Baidu's Institute of Deep Learning or even a Chinese research lab tied to the government or the People's Liberation Army, where open publication of such a breakthrough wouldn't be guaranteed. In some versions of that alternate timeline, ones in which researchers are pushed from an open U.S. lab to a closed Chinese one, the Chinese state could have gained exclusive access to one of the most powerful new methods in machine learning. It is highly unlikely that advance would have remained locked up in China permanently, but the possibility of the Chinese state getting years of head start applying such an advance should give pause to the notion that shutting down all U.S. labs in China is in U.S. interest.

Rebalancing AI R&D collaboration with China: A risk-based approach

The combination of China's deep integration into global AI R&D networks with the growing ethical and national security concerns about China's development and deployment of AI presents the U.S. and FCAI participants with the difficult task of rebalancing these networks. Some long-standing forms of AI R&D with Chinese partners—whether through collaborative research or hardware supplier relationships—now appear politically, ethically, or legally untenable. The conventional wisdom within parts of the technical community that scientific AI research should have no borders must be re-examined and adapted to consequential shifts in the power and potential applications of the underlying technology. We use the term "rebalancing" to express an approach that does not simply cut out China from these AI R&D networks altogether but that recognizes that business as usual with China is unsustainable.

Nevertheless, the risks from AI R&D with Chinese researchers should not lead to a wholesale severing of AI R&D networks between Chinese researchers and the U.S. and other democratic countries. This argument against a total decoupling of AI research isn't rooted in vague appeals to the win-win nature of all scientific collaboration. Instead, it derives from an assessment of the current realities of international AI R&D networks, how AI knowledge disseminates. and the ongoing benefits for the U.S.

Balancing these competing considerations in AI R&D with China should not be left entirely to the decisions of individual scientists, universities, or corporations. Their decisions require careful consideration of the risks involved—political, ethical, and technological—and regular re-examination informed by government sharing of information regarding risks and threats and, in some cases, carefully targeted changes to regulatory boundaries.

To achieve such a rebalancing will require the U.S. and like-minded partners to develop a risk-based approach that accounts for the current realities of both the technology and the geopolitics, while preserving benefits of international R&D networks, advancing progress on the goals outlined. To do so requires particularized accounting of the full range of costs and benefits in each of the pathways of AI R&D collaboration discussed in Part II. This part discusses more fully why rebalancing AI R&D networks with China is necessary and outlines a framework for developing a risk-based approach for each of these

pathways, including a range of ways that the government can inform the risk assessment process as well as what researchers and institutions individually and collectively can do to implement this risk-based approach.

This rebalancing will not be effective if it is entirely unilateral because of the global nature of AI R&D and China's significant role not only in U.S. research ecosystems but also those of allies that are leaders in AI, including those participating FCAI. International cooperation and alignment will therefore be essential. Otherwise, any U.S. action to reduce risks from AI knowledge transfers to Chinese researchers will require international cooperation to avoid leakage—leading Chinese researchers and students to seek AI knowledge, collaborations, and learning elsewhere. Such an outcome would both undermine the U.S. response and degrade U.S. research capacity. For instance, complete bans on Chinese students studying at U.S. universities or on collaboration on AI publications are unlikely to be followed by other governments.

Rebalancing collaboration with Chinese AI researchers should also include expanding international cooperation on AI among FCAI members and other AI R&D leaders. Indeed, the FCAI progress report highlighted the importance of cooperation in developing a risk-based approach to AI regulation, and this is another area where cooperation on a risk-based approach is needed.⁷⁵ The U.S. and FCAI partners are deepening cooperation on AI. For example, the 2019 update to the 2016 National AI R&D Strategic Plan recognizes the importance of international cooperation,⁷⁶ and a joint statement to leaders from the United States' Director of the White House Office of Science & Technology Policy and Canada's Minister of Innovation, Science, and Industry,⁷⁷ emphasized that "fostering stronger R&D engagement between academia and industry with our trusted allies is of mutual benefit to our economies and societies," noting the importance of working together in forums such as the G7 and the OECD. As outlined in a speech by National Security Advisor Jake Sullivan, the U.S. is strengthening international cooperation on digital technology across a range of forums in addition to the G7, such as the U.S.-EU Trade and Technology Council; developing digital trade rules as part of the Indo-Pacific Economic Framework; working with Australia, Japan, and India in the Quad; and has developed bilaterally on technology cooperation with South Korea.⁷⁸ Going forward, the U.S. should build on these approaches to include developing a risk-based approach to assessing ways to reduce R&D cooperation with China and to extend international cooperation with countries outside of the developed world but with large potential for building AI talent (e.g., India).

Such international cooperation should build on each country's experience in managing risks from research that can have national security or ethical

impacts. This has been true, for instance, of biomedical research for a while. The principles and procedures developed in the domestic context can inform a cooperative risk-based approach to AI R&D with China.

Applying a risk-based approach to AI knowledge transfers and collaboration with China on AI R&D

Adopting a risk-based approach to international AI R&D draws on wellestablished methodologies to assess risks and weigh costs and benefits in order to identify appropriate courses of action.⁷⁹ The key departure here is the application of risk and cost-benefit analysis in a way that expands the aperture to respond to the changes outlined above in the geostrategic environment with respect to China and the role of AI in this context. Given these developments, there is a role for governments to provide the information needed to inform the assessment of risks, costs, and benefits in AI R&D with China, to help coordinate across researchers, universities and labs, and to support international cooperation and alignment.

The range of actors and equities at stake in any reassessment of U.S.-China AI R&D networks will require developing a multistakeholder approach that is capable of understanding different viewpoints as well as the different costs/benefits that arise. For instance, research institutions and researchers all need timely and accurate information about prospective Chinese researchers in order to assess the risks from R&D collaboration on AI. Given that China's AI capability presents the U.S. and its allies with a range of economic and security risks, any cost-benefit analysis will need to take into account not just the private costs or benefits, but also these broader interests. Indeed, the mismatch between how a researcher, university, or lab might assess costs/benefits, and the broader economic and national security costs/benefits to the U.S. from strengthening China's AI capacity points to a range of roles for the government. This could include providing the additional information needed for cost-benefit analysis to assess the broader social costs from collaboration and, where the costs are always too high to justify ongoing engagement, banning some forms of knowledge transfer and cooperation. For example, there might be a case to ban joint publication or enrollment in doctoral programs by researchers from any of China's "seven sons" that are linked to the CCP national security state.

In the remainder of this paper, we lay out how a risk-based framework for rebalancing cooperation with China should address each of the pathways of AI R&D collaboration discussed in Part II. Each of these operates differently within R&D ecosystems and presents different costs and benefits so that each warrants a different approach.

Enrollment of Chinese students

As shown in the data on overseas students discussed above, a significant amount of U.S.-China collaboration occurs in the U.S. as Chinese researchers come to the U.S. to study, with a majority choosing to remain in the U.S. Indeed, the broader importance of international students for U.S. higher education was outlined in a 2021 joint statement by the Department of State and Department of Education titled "A Renewed U.S. Commitment to International Education," which addressed the importance of international education as a benefit for U.S. national security by creating people-to-people links and improving U.S. innovation and competitiveness.⁸⁰

This is a core U.S. strength—the ability to attract and then retain top talent. Moreover, China's growing clampdown on academic freedom and access to information in China may increase the attractiveness of the U.S. and the West for many Chinese AI researchers.⁸¹ These benefits for the U.S. point to the importance of visa reform that increases the opportunities for international AI researchers-from China and elsewhere-to work in the U.S. after graduating.⁸² As Jake Sullivan recently noted, the U.S. needs to ensure that top talent can come and stay in the U.S.⁸³ Retaining risk-based calibrated opportunities for Chinese researchers to study in the U.S. dovetails with the broader debate in the U.S. about the need for visa reform to help attract international talent for research and work in the U.S. As recent research has underscored, immigrants are a key driver of U.S. innovation, accounting for about a guarter of U.S. startups and patents each year.⁸⁴ Indeed, one analysis of the costs to the U.S. of the reduction in Chinese students as a result of PP10043 is over \$210 billion over 10 years once account is taken of the expected value of lost patents in addition to lost tuition.⁸⁵

Going forward, there are a range of steps the U.S. should take to expand opportunities for AI talent to study and work in the U.S. These include efforts to streamline the immigration process and reduce the barriers for those with advanced degrees in STEM fields to work and change employers in the U.S. For Chinese students pursuing advanced degrees in STEM fields, the U.S. should restructure visa processing to frontload the uncertainty for applicants. Following a 2018 order from the Trump administration,⁸⁶ Chinese students in certain STEM fields must reapply for their visas each year if they exit the country. The long delays and annual uncertainty of that process make it extremely risky for Chinese students to participate in doctoral programs in the U.S., knowing that they could lose access to the country any one of the four to seven years such a program can take. Given the volume of applications each year and the limited tools visa officers have to learn about applicants, repeating this exercise each year likely provides little additional insight for the U.S. while adding great uncertainty for the applicants. The process would be better served by a thorough upfront screening of applicants—one that would give visa officers the time and resources they require—and if they are approved, the visas they receive have longer validity.

The U.S. should also work with other governments to support their own approaches to assessing costs/benefits from access of Chinese students to AI programs and collaborative AI research. This can include common principles to guide the cost-benefit approach as well as supporting information exchange between U.S. law enforcement and other national security agencies and their partner agencies in these other countries with the aim of improving the cost-benefit analysis.

Collaborative research with Chinese partners

The highly diffuse and international AI research landscape means that many decisions about cooperation on specific research projects fall to individual researchers and institutions. But governments do have a role to play in that process by providing researchers and institutions with the information and support to make responsible

We propose that these decisions be grounded in three main considerations: (1) access to knowledge and resources; (2) the background of the Chinese partners; and (3) likely applications of the research. Here we present these considerations in the form of questions that researchers or institutions should ask themselves before embarking on such a project:

- Is this cooperative project giving the Chinese partners access to new resources—novel engineering knowhow, nonpublic datasets, etc.—that they would not have on their own and that they would not be able to acquire from simply reading the publicly released research paper?
- 2. Do you know who those Chinese partners are and whether they or their institution have ties to the Chinese military or security apparatus, in particular the Seven Sons of National Defense?
- 3. Have you thought through the likely applications of this research, specifically in China, and are you comfortable with the work contributing to those applications?

Of these three questions, the first should be largely answerable by researchers and their institutions. They are in the weeds of the research itself and know how it will be conducted between the two institutions. For the second question, governments can have an important role to play. And answering the third question will often involve consultation between these parties.

The first question is rooted in the reality of AI knowledge diffusion described in the above sections. As discussed above, for the majority of AI research conducted today, almost the entire value of the research will be contained in the publicly available research paper that will be posted online. In these cases, the cooperative nature of the project likely does not give the Chinese partners access to novel resources or knowledge they would not have in the absence of cooperation. In these cases, the answer to the first question is likely "no."

But there are types of research—large model construction, or hardware-driven areas—in which there may be meaningful knowledge transferred in the act of conducting the research. In addition, some joint research may provide the Chinese partner access to additional resources, such as unique datasets, that they otherwise would not have access to. In these cases, the answer to the first question would be "yes."

A "yes" answer to this first question is not by itself a reason to avoid research cooperation. In some cases, such as large-scale climate modeling or other projects with clear public benefits, sharing compute, data, or engineering knowhow may in fact be justifiable. In all cases, it will require an assessment of costs and benefits. This then leads to the second question: Do you know who your partner is?

Chinese public and private institutions often have complex and difficult to define relationships with parts of the Chinese government and the Communist Party. Not all of the relationships are particularly close, and not everything done within a company or research institution is done at the behest of the state. The Party and state simply do not have the expertise, resources, or desire to dig into the operational details of every Chinese institution. But given the Chinese state's heavy focus on AI and the many unethical uses to which it applies the technology, an AI researcher cooperating with a Chinese partner should understand their partner's relationship to the Chinese state.

In terms of capabilities, ethics, and other risks stemming from AI R&D, the two most concerning aspects of the Chinese state are its military and its maze of intelligence and internal security organs. Even with this limited scope, the blurry lines created by the PRC's civil-military fusion (CMF) initiative make bright lines difficult. In the most expansive interpretation put forward by some international analysts, the very existence of CMF makes almost every actor in

Given the Chinese state's heavy focus on Al and the many unethical uses to which it applies the technology, an Al researcher cooperating with a Chinese partner should understand their partner's relationship to the Chinese state. China a participant in the country's military buildup. Proponents of this view, point to Article 77 of China's 2015 National Security Law, which requires citizens and organizations to "provide necessary support" to China's public security and military organs. Other scholars view CMF as China's attempt to recreate aspects of the U.S.' vaunted military-industrial complex, with an added layer of coercive power available to the PRC state. It is beyond the scope of this paper to adjudicate this debate in general or as it applies in the context of R&D collaboration.

As a minimum baseline, however, we suggest institutions at least review research partnerships with the so-called Seven Sons of National Defense, a group of PRC universities with strong ties to the country's defense industry. According to one study, three quarters of graduates recruited by PRC state-owned defense companies came from the Sevens Sons.⁸⁷ While not every researcher affiliated with these universities is themselves affiliated with the Chinese military, partnerships here should receive extensive scrutiny. In proposing a AI research partnership with the "Seven Sons," there could either be a complete ban on AI R&D cooperation or additional scrutiny could apply when answering the first question—are you giving partners new capabilities they wouldn't otherwise have? But fully understanding one's partner will often necessitate the assistance of governments.

To facilitate this assistance, we suggest establishing a new public-private research security institution. This institution would be a place in which academia, industry, civil society, and government can share information and best practices on protecting research integrity from all threats, including those originating in the PRC. One proposal for such an institution came from a policy brief from the Center for Security and Emerging Technology.⁸⁸ A version of this proposal was included in the Senate's long-debated U.S. Innovation and Competition Act,⁸⁹ but it was eventually dropped from the final Chips and Science bill passed by both chambers.

As proposed, the new research security institution would be an industry- and academia-led body that can help inform those confronting either ethically or geopolitically difficult questions about research security in relation to international cooperation. It would provide a venue for many public and private sector actors to share information and engage with the U.S. government outside of the highly fraught circumstances of punitive law enforcement actions. The new institution would receive initial support and ongoing participation by government agencies, but it would be led by a coalition of groups from the research community itself: universities, companies, and labs. This will be key to assuaging some of the mutual suspicions that arise when government actors attempt to dictate terms or punish participants in the research community. While existing proposals for this institution have centered on the United States, this model could be applied within and across democratic countries in different ways. Each country could seek to create its own research security institution, one that accounts for the nuanced needs and constraints of their own domestic environment. Or a small coalition of countries—perhaps beginning with Five Eyes—could seek to create a transnational research security institution, one that leveraged the experience of researchers across several countries to create the most comprehensive information-sharing platform for making security- and ethics-informed research decisions.

Whatever form such an institution takes, it will be crucial for governments to remember that their role is to support and contribute to—but not control—the decisions made by researchers. There already exist laws and regulations for punishing those who directly violate laws around intellectual property, sanctions, or export controls. In order to gain the needed buy-in from industry and the research community, the institution would need to be a cooperative endeavor in which all partners bring their unique knowledge and expertise to bear on shared issues of concern.

This role for government could be supplemented by more specific policies to help guide universities and research institutions, particularly when it comes to areas of AI research where learning presents heightened economic or security risks for the U.S. This could draw on previous work in areas such as cybersecurity and biotechnology. For example, the U.K. National Cyber Security Centre Trusted Research Guidance for Senior Leaders and the U.S. joint Biotechnology and Biological Sciences Research Council and Medical Research Council and Welcome Trust policy statement provides principles and guidelines aimed at helping research institutions and researchers manage the risks.

The FBI is already seeking to expand partnerships with universities and research institutions to help these organizations understand the threat and develop effective mitigation strategies.⁹⁰ This type of partnership could be expanded to include other government agencies that can also provide information, for instance, on how AI research can be used in ways that is dual-use and unethical in third countries such as China. Further, more detailed articulation of principles for trustworthy and ethical AI beyond the high-level principles that currently exist would be useful as researchers and institutions navigate how to balance AI R&D against potential unethical use of that AI.

The issue of managing risks from research while also sustaining an internationally open research environment has also received attention from the G7. Specifically, the 2021 G7 Research Compact established the G7

Whatever form such an institution takes, it will be crucial for governments to remember that their role is to support and contribute to—but not control—the decisions made by researchers. Working Group on the Security and Integrity of the Research Ecosystem, which has been tasked with developing principles that can help protect the research and innovation ecosystem across the G7 from risks to open and reciprocal research collaboration, while also preserving the principles of open science and research freedom and independence. Given the complexities and pathways for A& R&D, the G7 should seek to expand and target this work for specific technologies such as AI.

Open publication of AI research

Given the diversity of applications for any fundamental AI innovation, the research norms of open publication sometimes present thorny tradeoffs. Open publication means an incremental advance in network architecture made by a Swiss researcher can quickly be put to work improving a cancer diagnosis system in Pittsburgh. But it also may improve the performance of a Chinese facial recognition system targeting dissidents. Such troubling use cases make it tempting to look for ways to impose new controls on who can access online research repositories, but practical and cultural hurdles make such controls difficult to implement effectively.

On the practical front, while democratic governments could use administrative regulations to force companies to "geo-fence" these online repositories and block access to them from China, this would likely prove futile. The same tools used to circumvent the Great Firewall could be used to circumvent the geo-fence, or a single person outside of China with an internet connection and a hard drive could simply download the data and bring it back to China.

And on the cultural front, barriers may be even more difficult to achieve. Al researchers are a large and highly international community. A 2020 study estimated that there were over 86,000 researchers who uploaded AI papers to ArXiv that year,⁹¹ a number that likely tops 100,000 this year. And in one study of the authors at an elite AI conference, over half of them were living and working in a different country from where they completed their undergraduate degree.⁹² A 2021 policy brief on research security by CSET described the challenge as follows: "[M]any, if not most, American researchers are unfamiliar with law enforcement, skeptical of their motives, and wary of restrictions on scientific openness and collaboration. Because of this, researchers are often unwilling to proactively collaborate with the government, including but not limited to law enforcement". As the norms around open publication have been deeply woven into the culture of this community, attempts by governments to force a dramatic restriction in AI research dissemination are likely to encounter heavy resistance and even outright defiance.

The above realities—a strong norm of open publication, along with practical and cultural hurdles to altering that norm—strongly constrain the ability of governments to curb China's access to global AI research advances. Given the government's limited ability to alter this status quo and the likelihood of a backlash if it tries, we recommend devoting attention and resources to the other vectors of cooperation.

Al conferences

As outlined above, the papers and presentations of most AI conferences are public, so reducing access to the learning that attending AI conferences provide would require the types of steps that reduce access to online information more broadly, which as discussed above, is likely not feasible nor in the interest of the United States. Moreover, given important advances in China on AI, the learning and networking opportunities at conferences also can provide U.S. researchers with valuable insights. Besides screening for clearly bad actors, such as spies, attendance at these conferences should be kept open.

Private sector AI labs in China

The decision on whether private companies should establish or maintain Al labs in China needs to contend with a range of costs and benefits on both ethical and strategic dimensions. Ethically, these considerations center on whether the research conducted there will contribute to oppressive technology deployment in China. Strategically, benefits have historically included access to foreign technical talent that powers American businesses; strategic costs have traditionally been conceptualized as hands-on training and a transfer of knowledge from the more advanced U.S. company to the less technically advanced Chinese researchers.

But dramatic changes of the past decade have complicated, or even inverted, some of these questions. China's dramatic escalation of Orwellian surveillance has made the ethical concerns much more pressing, but the transition to fully open and online research publication has also meant that fundamental AI research can contribute to China's surveillance state regardless of whether it is done in Boston or Beijing. And China's rapid arrival on the global frontier in AI research means that, today, transfers of knowledge often flow in reverse, from the China-based lab back to the U.S. company and the global research ecosystem. There was a time when the U.S. could have significantly slowed the growth of China's AI capabilities by eliminating these overseas labs, but that time has likely passed. As illustrated by the case of ResNet, the growing risk is that if Chinese researchers are pushed out of

American labs and into Chinese state-directed labs, there is a chance that breakthroughs made there could be kept bottled up within China.

Given this balance of risks, we recommend that U.S. research labs be allowed to continue operating in China but that they operate within constraints meant to minimize the ethical and strategic concerns involved. Ethically, these labs should only conduct research that is intended for open publication, thereby preventing knowledge transfer through proprietary or nonpublic work. While this won't stop research advances from contributing to China's surveillance state, it would limit any additional transfers due to that research being done in China.

In addition, these labs should not forge partnerships or recruit from institutions with substantial military ties. Microsoft's Beijing lab has reportedly taken a step in this direction by halting its recruiting from the aforementioned Seven Sons of National Defense.⁹³ Finally, as with all foreign companies operating in China, these labs need to take major precautions to limit the odds that these labs are used as beachheads for intellectual property theft from the headquarters. These precautions can be very onerous for those forced to deal with them in their daily work, but they are worth the investment on both ethical and strategic dimensions. Indeed, new U.S. export controls adopted in October 2002, potentially expanding the list of entities subject to a ban or licensing on a range of foundational semiconductor technologies, will constrain the involvement of U.S. companies and individuals in transmitting this information or technology as well as the hiring of foreign researchers.⁹⁴

Conclusion: Bridging the government-researcher divide on Al cooperation

Following the above recommendations will require governments, private companies, and the research community to make concessions to the others' interests, expertise, and worldview. Governments may be uncomfortable acknowledging that they cannot simply dictate the terms of research cooperation or Al knowledge diffusion globally. Al researchers and their institutions may also feel uncomfortable engaging with the ethical and geopolitical dimensions of their work and unequipped to make those assessments on their own. But both camps must move past that discomfort and engage in good faith attempts to tackle these questions together. In parallel, the U.S. government will need to make concessions to the interests and desires of its partners and accept that a slightly more accommodating multilateral initiative will be far more effective than a hardline unilateral initiative.

As outlined above, to support research institutions in making informed decisions, governments should help establish a new public institution focused on protecting research integrity and security.

This working paper has also outlined steps that can be taken to address concerns about the growth in Chinese AI capabilities and unethical deployments of the technology, without shutting the door completely on collaboration and information exchange with China on AI. The above recommendations also underscore the importance of governments and research institutions working harder to develop a more geopolitically sustainable and ethically defensible equilibrium when it comes to AI research cooperation with China.

Appendix:

Research papers co-authored by researchers in FCAI countries and China, 2010-2021

United States:









Australia:





Canada:

Japan:





Singapore:

Source: Emerging Technology Observatory's Country Activity Tracker: Artificial Intelligence⁹⁵

References

- 1. Cameron F Kerry et al., "Strengthening International Cooperation on AI," n.d., 123.
- 2. Kerry et al., "Strengthening International Cooperation on AI," 123.
- 3. Autumn Toney and Melissa Flagg, "Research Impact, Research Output, and the Role of International Collaboration," *Center for Security and Emerging Technology* (blog), accessed November 22, 2022, https://cset.georgetown.edu/publication/research-impact-research-output-and-the-role-of-international-collaboration/.
- Center for Security and Emerging Technology, William Hannas, and Huey-Meei Chang, "China's Access to Foreign AI Technology" (Center for Security and Emerging Technology, September 2019), https://doi.org/10.51593/20190053.
- Evan Feigenbaum, "The Deep Roots and Long Branches of Chinese Technonationalism -Carnegie Endowment for International Peace," accessed November 22, 2022, https://carnegieendowment.org/2017/08/12/deep-roots-and-long-branches-of-chinesetechnonationalism-pub-72815.
- 6. Jeffrey Ding, "ChinAI #37: Happy 20th Anniversary to Microsoft Research Asia," accessed November 22, 2022, https://chinai.substack.com/p/chinai-37-happy-20th-anniversary.
- Gregory Huang, "The World's Hottest Computer Lab," MIT Technology Review, accessed November 22, 2022, https://www.technologyreview.com/2004/06/01/232827/theworlds-hottest-computer-lab/.
- 8. Matt Sheehan, "The Chinese Way of Innovation," April 26, 2022, https://www.foreignaffairs.com/articles/china/2022-04-21/chinese-way-innovation.
- 9. Matt Sheehan, "WFH: Chinese Engineers Abroad Come Back | WIRED," accessed November 22, 2022, https://www.wired.com/story/excerpt-transpacific-experimentmatt-sheehan/.
- Jordan Novet, "Why a Deep-Learning Genius Left Google & Joined Chinese Tech Shop Baidu (Interview)," VentureBeat (blog), July 30, 2014, https://venturebeat.com/business/andrew-ng-baidu/.
- 11. Graham Webster et al., "Full Translation: China's 'New Generation Artificial Intelligence Development Plan' (2017)," New America, accessed November 22, 2022, http://newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinasnew-generation-artificial-intelligence-development-plan-2017/.
- 12. Matt Sheehan, "How China's Massive AI Plan Actually Works," MacroPolo, accessed November 22, 2022, https://macropolo.org/analysis/how-chinas-massive-ai-planactually-works/.
- Jackie Snow, "China's AI Startups Scored More Funding than America's Last Year | MIT Technology Review," accessed November 22, 2022, https://www.technologyreview.com/2018/02/14/145616/chinas-ai-startups-scoredmore-funding-than-americas-last-year/.
- 14. Fei-Fei Li, "Opening the Google Al China Center," Google, December 13, 2017, https://blog.google/around-the-globe/google-asia/google-ai-china-center/.

- 15. Adam Conner-Simons, "CSAIL Launches New Five-Year Collaboration with IFlyTek," MIT News | Massachusetts Institute of Technology, accessed November 22, 2022, https://news.mit.edu/2018/csail-launches-five-year-collaboration-with-iflytek-0615.
- 16. China Institute for Science and Technology Policy, "China AI Development Report 2018" (Tsinghua University, July 2018).
- 17. "Global AI Vibrancy Tool," accessed November 22, 2022, https://aiindex.stanford.edu/vibrancy/.
- "The 10 Most Valuable Private AI Companies CB Insights Research," accessed November 22, 2022, https://www.cbinsights.com/research/most-valuable-private-aicompanies/.
- Ryan Mac, "The US Just Blacklisted China's Most Valuable Facial Recognition Startups Over Human Rights Abuses," accessed November 22, 2022, https://www.buzzfeednews.com/article/ryanmac/trump-us-blacklist-megvii-sensetimeyitu-xinjiang.
- 20. "CSET PARAT," accessed November 22, 2022, https://parat.cset.tech/.
- 21. Center for Strategic & International Studies, *China's Power: Up for Debate 2021 Debate 2*, 2021, https://www.youtube.com/watch?v=PiwxEpF4QbI.
- 22. Matt Sheehan, "China Technology Forecast in 2025: Fragile Tech Superpower," MacroPolo, accessed November 22, 2022, https://macropolo.org/analysis/chinatechnology-forecast-2025-fragile-tech-superpower/.
- 23. Center for Security and Emerging Technology et al., "AI Definitions Affect Policymaking" (Center for Security and Emerging Technology, June 2020), https://doi.org/10.51593/20200004.
- 24. Gleb Chuvpilo, "Who's Ahead in Al Research at NeurIPS 2020?," *Medium* (blog), May 20, 2022, https://chuvpilo.medium.com/whos-ahead-in-ai-research-at-neurips-2020-bf2a40a54325.
- 25. "The Global AI Talent Tracker," MacroPolo, accessed November 22, 2022, https://macropolo.org/digital-projects/the-global-ai-talent-tracker/.
- 26. "Global AI Talent Tracker," MacroPolo, https://macropolo.org/digital-projects/the-globalai-talent-tracker/.
- 27. Megha Rajagopalan, "This Is What A 21st-Century Police State Really Looks Like," BuzzFeed News, accessed November 22, 2022, https://www.buzzfeednews.com/article/meghara/the-police-state-of-the-future-isalready-here.
- 28. Paul Mozur, "One Month, 500,000 Face Scans: How China Is Using A.I. to Profile a Minority - The New York Times," accessed November 22, 2022, https://www.nytimes.com/2019/04/14/technology/china-surveillance-artificialintelligence-racial-profiling.html.
- 29. Mozur, "One Month, 500,000 Face Scans," https://www.nytimes.com/2019/04/14/technology/china-surveillance-artificialintelligence-racial-profiling.html.
- 30. Wang Yiwei, "Face Recognition Spots Wanted Suspects at Qingdao Beer Festival," Sixth Tone, August 31, 2017,

https://www.sixthtone.com/news/1000773/https%3A%2F%2Fwww.sixthtone.com%2Fn ews%2F1000773%2Fface-recognition-spots-wanted-suspects-at-qingdao-beer-festival.

31. Paul Mozur, "Let's Talk About China," *The New York Times*, September 21, 2018, sec. Reader Center, https://www.pytimes.com/interactive/2022/uri/embeddedinteractive/54110634-8e67

https://www.nytimes.com/interactive/2022/uri/embeddedinteractive/54110634-8e67-5c2f-99bd-af9b2542304f?

- 32. "A Quick Guide to the US-China Trade War BBC News," accessed November 22, 2022, https://www.bbc.com/news/business-45899310.
- 33. "Quick Guide to US-China Trade War" https://www.bbc.com/news/business-45899310.
- 34. Will Knight, "MIT Cuts Ties With a Chinese AI Firm Amid Human Rights Concerns," *Wired*, accessed November 22, 2022, https://www.wired.com/story/mit-cuts-ties-chinese-ai-firm-human-rights/.
- 35. Knight, "MIT Cuts Ties With Chinese AI Firm," https://www.wired.com/story/mit-cutsties-chinese-ai-firm-human-rights/.
- 36. Yang Jingmingmin, "MSRA被曝停招国防七子及北邮学生," accessed November 22, 2022, https://mp.weixin.qq.com/s/yKWVd9DSRy7nhjtkR7UTMw?fbclid=IwAR36yLtMZmevnd mEIJ-Vc_AZ5WImrID-SP5F0FttM_PcowIA4muKpo-Hlvk.
- 37. Edmund Andrews, "China and the United States: Unlikely Partners in AI," Stanford HAI, accessed November 22, 2022, https://hai.stanford.edu/news/china-and-united-states-unlikely-partners-ai.
- 38. "US Considered Ban on Student Visas for Chinese Nationals," *Financial Times*, October 2, 2018.
- 39. Jeffrey Mervis, "More Restrictive U.S. Policy on Chinese Graduate Student Visas Raises Alarm | Science | AAAS," accessed November 22, 2022, https://www.science.org/content/article/more-restrictive-us-policy-chinese-graduatestudent-visas-raises-alarm.
- 40. Ellen Nakashima, "Charges Dismissed against MIT Professor Accused of Hiding Research Ties to China," *Washington Post*, accessed November 22, 2022, https://www.washingtonpost.com/national-security/mit-gang-chendismiss/2022/01/20/912f68aa-786b-11ec-bf97-6eac6f77fba2_story.html.
- 41. Natasha Gilbert, "I Lost Two Years of My Life': US Scientist Falsely Accused of Hiding Ties to China Speaks Out," *Nature* 603, no. 7901 (March 7, 2022): 371–72, https://doi.org/10.1038/d41586-022-00528-2.
- 42. Gilbert, "'Lost Two Years of My Life'," 371–72, https://doi.org/10.1038/d41586-022-00528-2.
- 43. Raquel Leslie and Brian Liu, "U.S. Justice Department Puts an End to Controversial China Initiative," Lawfare, accessed November 22, 2022, https://www.lawfareblog.com/.
- 44. Juro Osawa and Shai Oster, "Beijing Tightens Grip on ByteDance by Quietly Taking Stake, China Board Seat," The Information, accessed November 22, 2022, https://www.theinformation.com/articles/beijing-tightens-grip-on-bytedance-by-quietlytaking-stake-china-board-seat.
- 45. International Students in Science and Engineering, National Foundation for American Policy, NFAP Policy Brief August 2021

- 46. "Higher Education in Science and Engineering | NSF National Science Foundation," accessed November 22, 2022, https://ncses.nsf.gov/pubs/nsb20197/international-s-e-higher-education.
- 47. "Global AI Talent Tracker," MacroPolo, https://macropolo.org/digital-projects/the-globalai-talent-tracker/.
- 48. Paul Mozur and Cade Metz, "A U.S. Secret Weapon in A.I.: Chinese Talent," The New York Times, June 9, 2020, https://www.nytimes.com/2020/06/09/technology/china-ai-research-education.html.
- 49. "Higher Education in Science and Engineering | NSF National Science Foundation," accessed November 22, 2022, https://ncses.nsf.gov/pubs/nsb20197/international-s-e-higher-education.
- 50. "Suspension of Entry as Nonimmigrants of Certain Students and Researchers From the People's Republic of China," Federal Register, June 4, 2020, https://www.federalregister.gov/documents/2020/06/04/2020-12217/suspension-of-entry-as-nonimmigrants-of-certain-students-and-researchers-from-the-peoples-republic.
- 51. Russell Flannery, "American Universities Are Losing Chinese Students To Rivals: U.S.-China Business Forum," Forbes, accessed November 23, 2022, https://www.forbes.com/sites/russellflannery/2022/08/11/american-universities-arelosing-chinese-students-to-rivals-us-china-business-forum/.
- 52. Autumn Toney and Melissa Flagg, "Research Impact, Research Output, and the Role of International Collaboration," *Center for Security and Emerging Technology* (blog), accessed November 22, 2022, https://cset.georgetown.edu/publication/research-impact-research-output-and-the-role-of-international-collaboration/.
- 53. Husanjot Chahal et al., "Country Activity Tracker," *Center for Security and Emerging Technology* (blog), accessed November 23, 2022, https://cset.georgetown.edu/publication/country-activity-tracker/.
- 54. Chahal et al., "Country Activity Tracker," https://cset.georgetown.edu/publication/country-activity-tracker/.
- 55. Chahal et al., "Country Activity Tracker," https://cset.georgetown.edu/publication/country-activity-tracker/.
- 56. Chahal et al., "Country Activity Tracker," https://cset.georgetown.edu/publication/country-activity-tracker/.
- 57. Chahal et al., "Country Activity Tracker," https://cset.georgetown.edu/publication/country-activity-tracker/.
- 58. "Al Index 2022," Stanford Institute for Human-Centered Artificial Intelligence, accessed November 23, 2022, https://hai.stanford.edu/research/ai-index-2022.
- 59. Dave Gershgorn, "The Data That Transformed AI Research—and Possibly the World," Quartz, July 26, 2017, https://qz.com/1034972/the-data-that-changed-the-direction-ofai-research-and-possibly-the-world/.
- 60. "Papers with Code The Latest in Machine Learning," accessed November 23, 2022, https://paperswithcode.com/.
- 61. "Advances in Neural Information Processing Systems Datasets and Benchmarks," n.d., https://papers.nips.cc/.

- 62. "NeurIPS 2022," accessed November 23, 2022, https://neurips.cc/Conferences/2021/PaperInformation/PaperChecklist.
- 63. Irene Solaiman et al., "Release Strategies and the Social Impacts of Language Models," n.d., 71.
- 64. Will Heaven, "OpenAl's New Language Generator GPT-3 Is Shockingly Good—and Completely Mindless," MIT Technology Review, accessed November 23, 2022, https://www.technologyreview.com/2020/07/20/1005454/openai-machine-learninglanguage-generator-gpt-3-nlp/.
- 65. Karen Hao, "The messy, secretive reality behind OpenAI's bid to save the world", MIT Technology Review. https://www.technologyreview.com/2020/02/17/844721/ai-openaimoonshot-elon-musk-sam-altman-greg-brockman-messy-secretive-reality/.
- 66. Will Heaven, "Meta Has Built a Massive New Language AI—and It's Giving It Away for Free," MIT Technology Review, accessed November 23, 2022, https://www.technologyreview.com/2022/05/03/1051691/meta-ai-large-languagemodel-gpt3-ethics-huggingface-transparency/.
- 67. Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models" (arXiv, July 12, 2022), https://doi.org/10.48550/arXiv.2108.07258.
- 68. "Al Index 2022," Stanford Institute for Human-Centered Artificial Intelligence, accessed November 23, 2022, https://hai.stanford.edu/research/ai-index-2022.
- 69. Sergei Ivanov, "NeurIPS 2020. Comprehensive Analysis of Authors, Organizations, and Countries.," *Criteo R&D Blog* (blog), October 15, 2020, https://medium.com/criteo-engineering/neurips-2020-comprehensive-analysis-of-authors-organizations-and-countries-a1b55a08132e.
- 70. Ivanov, "ICML 2020. Comprehensive Analysis," https://medium.com/criteoengineering/icml-2020-comprehensive-analysis-of-authors-organizations-and-countriesc4d1bb847fde.
- 71. Roxanne Heston and Remco Zwetsloot, "Mapping U.S. Multinationals' Global AI R&D Activity, CSET Issues Brief," December 2020.
- 72. Heston and Zwetsloot, "Mapping U.S. Multinationals", December 2020.
- 73. Klon Kitchen and Bill Drexel, Pull US AI Research Out of China, Ideas, August 10, 2021.
- 74. Heston and Zwetsloot, "Mapping U.S. Multinationals", December 2020.
- 75. Cameron F Kerry, Joshua P Meltzer, Andrea Renda, Alex C Engler, and Rosanna Fanni. "Strengthening International Cooperation on Al," n.d., 123.
- 76. "The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update," 2019, 50.
- 77. Eric Lander, "Joint Statement to Leaders from the United States' Director of the White House Office of Science & Technology Policy and Canada's Minister of Innovation, Science and Industry | OSTP," The White House, accessed November 23, 2022, https://www.whitehouse.gov/ostp/news-updates/2021/11/18/joint-statement-toleaders-from-the-united-states-director-of-the-white-house-office-of-science-technologypolicy-and-canadas-minister-of-innovation-science-and-industry-2/.
- 78. Remarks by National Security Advisor Jake Sullivan at the Special Competitive Studies Project Global Emerging Technologies Summit, September 16, 2022.

- 79. Baruch Fischhoff, "The realities of risk-cost-benefit analysis" Science, Vol 360, No 6260, 30 Oct 2015.
- 80. "A Renewed U.S. Commitment to International Education" (Department of Education, July 27, 2021),
 - https://educationusa.state.gov/sites/default/files/intl_ed_joint_statement.pdf.
- 81. Emilie Tran, "China: Zero Tolerance for Academic Freedom," The Conversation, accessed November 23, 2022, http://theconversation.com/china-zero-tolerance-for-academicfreedom-85200.; Jennifer Ruth and Yu Xiao, "Academic Freedom and China," AAUP, October 10, 2019, https://www.aaup.org/article/academic-freedom-and-china.
- 82. America has an innovation problem. The H-1B visa backlog is making it worse, Vox, July 31, 2022 https://www.vox.com/future-perfect/23177446/immigrants-tech-companies-united-states-innovation-h1b-visas-immigration.
- 83. Remarks by National Security Advisor Jake Sullivan at the Special Competitive Studies Project Global Emerging Technologies Summit, September 16, 2022.
- 84. William R. Kerr and Sari Pekkala Kerr, Immigration Policy Levers for U.S. Innovation and Startups, Harvard Business School Working Paper 20-105.
- 85. International Students in Science and Engineering, National Foundation for American Policy, NFAP Policy Brief August 2022.
- 86. Josh Lederman and Ted Bridis, "AP Sources: US to Impose Limits on Some Chinese Visas | AP News," May 29, 2018,

https://apnews.com/article/82a98fecee074bfb83731760bfbce515.

- Center for Security and Emerging Technology, Ryan Fedasiuk, and Emily Weinstein, "Universities and the Chinese Defense Technology Workforce" (Center for Security and Emerging Technology, December 2020), https://doi.org/10.51593/20200043.
- 88. Melissa Flagg and Zachary Arnold, "A New Institutional Approach to Research Security in the United States," *Center for Security and Emerging Technology* (blog), accessed November 23, 2022, https://cset.georgetown.edu/publication/a-new-institutional-approach-to-research-security-in-the-united-states/.
- 89. Charles E. Schumer, "Text S.1260 117th Congress (2021-2022): United States Innovation and Competition Act of 2021," legislation, June 8, 2021, 2021/2022, http://www.congress.gov/.
- 90. David Bowdich, The Importance of Partnerships in Responding to the Chinese Economic Espionage Threat to Academia, Academic Security and Country Exploitation Annual Seminar, March 4, 2020.
- 91. "Global AI Talent Report 2020," accessed November 23, 2022, https://jfgagne.ai/globalai-talent-report-2020.
- 92. "Global AI Talent Tracker," MacroPolo, https://macropolo.org/digital-projects/the-globalai-talent-tracker/.
- 93. Yang Jingmingmin, "MSRA被曝停招国防七子及北邮学生," accessed November 22, 2022, https://mp.weixin.qq.com/s/yKWVd9DSRy7nhjtkR7UTMw?fbclid=IwAR36yLtMZmevnd mEIJ-Vc_AZ5WImrID-SP5F0FttM_PcowIA4muKpo-Hlvk.
- 94. "Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification," Federal Register, October 13, 2022,

https://www.federalregister.gov/documents/2022/10/13/2022-21658/implementation-of-additional-export-controls-certain-advanced-computing-and-semiconductor.

95. Husanjot Chahal et al., "Country Activity Tracker," *Center for Security and Emerging Technology* (blog), accessed November 23, 2022, https://cset.georgetown.edu/publication/country-activity-tracker/.

BROOKINGS

1775 Massachusetts Ave NW, Washington, DC 20036 (202) 797-6000 www.brookings.edu