

*Testimony before the U.S.-China Economic and Security Review Commission
“Technology, Trade, and Military-Civil Fusion: China’s Pursuit of Artificial Intelligence, New Materials, and New
Energy”*

by

*Helen Toner
Director of Strategy, Center for Security and Emerging Technology
Walsh School of Foreign Service, Georgetown University*

June 7, 2019

Madam Chairman, Madam Vice Chairman, members of the Commission: Many thanks for this opportunity to testify on this crucial and timely topic.

I am the Director of Strategy at the Center for Security and Emerging Technology. CSET is a new center at Georgetown University that was founded this past January to examine the security implications of new tech developments, including the kinds of questions addressed by this hearing. Our initial suite of research products—which we expect to begin making public this summer—combines our team’s expertise on China, national security and artificial intelligence to produce in-depth analysis of key issues related to U.S. competition with China on AI.

I previously worked on similar issues of AI and national security in San Francisco, including substantial engagement with AI researchers at corporate and university research labs in Silicon Valley and elsewhere. In order to deepen my understanding of the equivalent ecosystem in China, I spent most of 2018 living in Beijing, undertaking independent research and study.

In order to address the themes suggested by the Commission today, I’ll start with some scene-setting to describe what I see as key features of AI as a technology and address some common misconceptions. In brief, I’ll describe why AI should not be thought of as a typical dual-use technology; why AI research is characterized by an unusually open and collaborative environment, and what value that has; the potential costs to the United States of any restrictive measures that are not highly targeted; the role of human capital in AI research, and importance of strategic immigration policy to bolster U.S. competitiveness; China’s approach to data privacy, and why I believe discussions of human rights abuses in Xinjiang should not be closely tied to discussions of AI; and the current state of standardization efforts for AI. I’ll close with recommendations for how to strengthen the U.S. competitive advantage in AI.

1. AI as a general purpose technology

Artificial intelligence is a general purpose technology. The concept of a GPT, which comes from economics, refers to a technology that has the potential to significantly affect all sectors of society and the economy.¹

¹ Jovanovic & Rousseau (2005), *General Purpose Technologies*, <https://www.nber.org/papers/w11093>

Despite some overlap with the concept of “dual use,” which generally refers to militarily-useful technologies that also have some civilian applications, the terms are not synonyms. Unlike those of a dual-use technology, the military applications of a GPT generally represent only a small part of the technology’s overall usage and potential for value creation, rather than being one of its major facets. Commonly-cited examples of GPTs include electricity and information technology, in contrast with typical dual-use examples such as nuclear energy, rocketry, or biotechnology.

In the case of AI, we are already seeing promising applications across sectors with clear humanitarian implications, such as scientific innovation, healthcare, energy and transportation. Additionally, advances in technologies like speech recognition, translation, natural language processing and image processing can be applied across all sectors of the economy, spurring growth and making possible further new technologies and ways of living that are hard to imagine now. Consider how electricity not only led to artificial light, but also made possible elevators (and thereby high-rise buildings), telecommunications, modern agriculture, and myriad other inventions that revolutionized society. AI holds the promise of unleashing a similar transformation, and this has important implications for how governments should interact with this technology.

2. AI research norms and collaborations

a. Publishing norms

Since well before the beginning of the current boom in AI in around 2012, the field has been characterized by strong norms of open publishing. The vast majority of research progress is published on arXiv.org, a freely accessible repository for scientific papers maintained by Cornell University. These norms of openness are so strong that most major technology companies with AI research labs, including Google, Facebook, Amazon and Microsoft, also allow researchers to publish much of their work freely.

This open, distributed environment accelerates research progress in several ways. Researchers in one lab can easily test and build off of results published by a different lab; researchers in different labs (and even different countries) can straightforwardly collaborate on projects; researchers moving between jobs need less time to get settled with their new organization’s research and practices; less experienced researchers can easily teach themselves from online resources and quickly get to the level where they can contribute their own insights; and so on.

Because of AI’s substantial potential to benefit humanity, as described in the previous section, the boost to research progress provided by this structure is extremely valuable.

b. Research collaborations

Collaborations of many different kinds are a natural consequence of this model. Cooperative arrangements between universities and corporations are commonplace. Because many American companies with large AI research efforts (such as Google and Microsoft) operate around the globe, these companies have a wide range of partnerships with other entities in foreign countries, including China.

Overwhelmingly, these collaborations relate to very basic research that could have many potential applications. For example, Microsoft Research Asia—a Beijing-based research group that is one of China’s best AI labs, and certainly the most prestigious Western lab in China—announced a set of 40

research collaboration grants in December 2018, 23 of which went to Chinese institutions. Of these grants to Chinese institutions (listed alphabetically by author name), the first five topics relate to using AI for rehabilitation, improving conversational question-answering, segmenting objects in video footage, machine translation, and system architecture.² In other words, the typical US-China research collaboration represents a marginal improvement to a basic machine learning problem, contributing to the global commons of scientific research progress.

The most infamous example of a government-industry partnership is of course Google's work on, then withdrawal from, the U.S. Department of Defense's Project Maven. This withdrawal gave rise to a narrative sometimes heard in Washington that Silicon Valley refuses to work with the U.S. government, but is happy to cooperate with China. Contrary to this narrative, there are in fact many examples of arrangements in which U.S. tech companies (including Google) work with U.S. government partners.³ While the Washington-Silicon Valley relationship can be fraught and requires attention, oversimplifications of this kind should be avoided.

c. The costs of a restrictive approach

The openness of the current AI ecosystem can seem undesirable to policymakers concerned with shoring up the United States' technological advantage and securing American innovation, especially given AI's potential military applications. A natural impulse is to seek ways to close off external access to U.S. research, perhaps drawing inspiration from case studies like nuclear energy or rocketry.

However, an approach like this is likely to be counterproductive, given the general purpose nature of AI as described above and the field's interconnected global research environment. Because sensitive applications of AI represent such a small chunk of its potential uses, and non-sensitive applications hold such promise for promoting growth and prosperity, measures that attempt to broadly restrict access to AI research (for example, applications of export controls or restrictions on collaborative research that are not highly targeted) are likely to backfire in two mutually-reinforcing ways.

First, measures that restrict collaboration or open sharing of research are likely to slow down the pace of research progress within U.S. university and corporate labs, which would damage their standing on the world stage and reduce their market share in AI-enabled products and platforms.

Second, because AI workers are highly mobile, any such measures enacted in the United States are likely to prompt researchers to emigrate to other countries to continue their work unencumbered. Because AI is not primarily a military technology, even patriotic American experts may not see the case for staying in the United States if they will be more able to push the scientific frontier elsewhere. As described in the section on human capital below, researchers will quickly find employment overseas, including in the growing AI sectors of Canada and the UK (which are actively taking measures to

² Microsoft Research Lab – Asia (2018), *MSRA Collaborative Research 2019 Grant Awards Announcement*, <https://www.microsoft.com/en-us/research/lab/microsoft-research-asia/articles/msra-collaborative-research-2019-grant-awards-announcement/>

³ Examples include Google's work on DARPA programs on deepfakes and semiconductor design (<https://www.c4isrnet.com/it-networks/2019/03/13/forget-project-maven-here-are-a-couple-other-dod-projects-google-is-working-on/>), bids by Amazon and Microsoft on the Pentagon's cloud contract (<https://www.nytimes.com/2019/04/10/technology/amazon-microsoft-jedi-pentagon.html>), and Google and Facebook assisting the Census Bureau to defend against disinformation (<https://www.reuters.com/article/us-usa-census-fakenews-exclusive/exclusive-fearful-of-fake-news-blitz-u-s-census-enlists-help-of-tech-giants-idUSKCN1R812S>).

recruit AI talent). Not to mention the many Chinese and Russian researchers currently contributing their talents to the United States, who might be prompted to return to their home countries instead.

In short, I fear that attempts to bolster American competitiveness using restrictive measures will instead degrade U.S. leadership in science and technology, both due to the direct effect on U.S. research progress and due to the indirect effects of deterring talented workers from settling in the United States.

d. A framework for controls

In her recent testimony to the House Foreign Affairs Committee, New America Fellow Samm Sacks provides a useful framework for thinking about where to apply controls on technology:⁴

“In general, a technology should be subject to greater control if:

1. It is essential to military technology; however, the term “essential” should not be interpreted to encompass technology that is simply used or is usable by the military, since the defense industry is increasingly reliant on commercial off-the-shelf technology. The International Traffic in Arms Regulations (ITARs) are designed to fulfill this purpose, but differentiating between essential military technology (often controlled by the United States Munitions List) and dual-use technology remains a challenge;
2. There is a scarcity of knowledge about the technology, except among a small group of experts located in the United States or like-minded countries; and
3. The United States is truly ahead of the curve, and that technology is developed exclusively in the United States or other countries that enforce similar export controls. Technical experts must be regularly consulted to evaluate incremental differences between our technology and that of other countries on this point. Not doing so risks the “designing out” of U.S.-made components from products for global markets, which would advantage foreign companies with similar products that are not subject to export controls.”

This framework was provided in the context of export controls, but it applies equally well to international research collaborations. I urge the Commission to adopt this framework when considering what types of AI research might merit restriction, given that the vast majority of AI research does not meet any of these three criteria.

3. Human capital as a driver of AI progress

Access to skilled researchers and engineers is a key area of competition in the field of AI. The United States’ unique ability to attract and retain foreign talent represents, therefore, a key American advantage. See for example Figure 1 below, which shows the unique U.S. position as a massive net importer of patent holders.

China is working hard to catch up, with government initiatives like the Thousand Talents Plan (千人计划) and educational programs described in the April 2018 *Artificial Intelligence Innovation Action Plan*

⁴ Sacks (2019), *Samm Sacks Testifies Before House Foreign Affairs Committee on 'Smart Competition' With China*, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/samm-sacks-testifies-house-foreign-affairs-committee-smart-competition-china/>

[for] *Institutions of Higher Learning* (高等学校人工智能创新行动计划) aiming to step up the training of indigenous talent, and to incentivize Chinese abroad to return home.

The current U.S. approach stands in stark contrast to measures like this from China and other countries. Far from actively stepping up efforts to draw top foreign talent to its shores, recent changes in the U.S. immigration environment are actively undermining this historical—and critical—strength.

More than half of the computer scientists with graduate degrees working in the United States were born abroad (59% of workers with PhDs, 51% of those with Master's).⁵ Many of these workers came to the United States as international students, who disproportionately prefer to stay in the United States after completing their studies. More than 85% of Chinese and Indian students in U.S. computer science and engineering PhD programs state that they intend to stay after graduation, and the actual stay rates over five and ten years suggest that nearly this number do in fact remain.⁶

This status quo reflects the high quality of the U.S. commercial and research environment, as well as the attractiveness of the liberty, openness and prosperity found here.

A strategic approach to U.S. AI policy would seek to leverage these strengths in order to cement the United States' place as the standout global hub for AI talent. Unfortunately, recent changes in the immigration environment seem to be pushing in the opposite direction, with Executive Branch policies to constrain legal immigration to this country compounded by increasing reports of long delays in processing of Chinese students in STEM programs.

Ultimately, U.S. action or inaction that restricts the inflow of top-tier research talent from China is a dream come true for the Chinese government. Such actions do more than any Thousand Talents Plan ever could to bolster Chinese competitiveness.

4. Data as a driver of AI progress

a. Data privacy in China

Despite perceptions to the contrary, awareness of and concern about data privacy is rising among Chinese consumers, and the Chinese government is actively developing laws and regulations in response. This push is part of the country's larger effort to build a complex governance regime for cyberspace and information and communications technology.⁷

In broad strokes, this data privacy regime seeks to protect consumer privacy from technology companies working in China. A standard called the Personal Information Security Specification (个人信息安全规范) took effect in May 2018 and forms the first piece of the regime.⁸ Although

⁵ National Science Foundation (2015), *Figure 3-32: Foreign-Born Scientists and Engineers Employed in S&E Occupations, by Highest Degree Level and Broad S&E Occupational Category: 2015*, <https://www.nsf.gov/statistics/2018/nsb20181/assets/901/figures/fig03-32.pdf>

⁶ National Science Foundation (2015), *Appendix Table 3-21: Plans of Foreign Recipients of U.S. Doctorates to Stay in the United States, by Field of Doctorate and Place of Origin: 2004–15*, <https://www.nsf.gov/statistics/2018/nsb20181/assets/901/tables/at03-21.pdf>.

⁷ Sacks, *China's Emerging Cyber Governance System*, <https://www.csis.org/chinas-emerging-cyber-governance-system>

⁸ Sacks, Shi, & Webster (2019), *The Evolution of China's Data Governance Regime: A Timeline*, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/china-data-governance-regime-timeline>

modeled heavily on the European Union's General Data Protection Regulation (GDPR), this specification seeks to be somewhat more permissive than GDPR in order to be more business friendly.⁹

One consideration for the United States is whether China will be able to significantly affect global data privacy practices simply by virtue of beginning to regulate companies operating in China before the United States regulates companies operating here. GDPR is already considered to have played a significant role in setting the parameters for future privacy conversations, because multinational companies that operate in Europe have had to build out compliance structures based on the European law, and will likely use those same structures to implement any future legislation.

b. Data as a strategic resource

One related note worth delving into is that I believe the idea of data as a general-purpose strategic resource ("the new oil") has been exaggerated. While it is true that data is an important input to AI systems, data is not generically useful for training any kind of system. This is because AI systems are essentially pattern-recognition machines. Any given AI application will require data that is relevant to the specific problem it is trying to solve, from which it can learn what kinds of patterns are likely to exist in similar data. For example, data on consumers' purchasing history is valuable for predicting future purchasing behavior, data collected by autonomous vehicles can be used to improve autonomous vehicle algorithms, and so on.

In other words, even if China's laws and norms around consumer data privacy remained significantly laxer than the United States', that would not necessarily have many implications beyond the possibility that Chinese companies could more effectively sell their products to Chinese consumers.

Inasmuch as it makes sense to think of data as conferring a strategic advantage, a more productive approach might be to identify specific applications of concern, consider what data would be required to train those systems, and work to improve U.S. access to that type of data. Notably, the United States appears well-positioned in several security-relevant domains, for example due to the fact that the United States has far more platforms and bases, in many more environments, collecting military-relevant data from many more sensors, than China.

Efforts to utilize and protect valuable existing datasets like this would be much more beneficial to the United States than worrying excessively about Chinese citizens' attitudes to privacy.¹⁰

c. The role of data in digital authoritarianism

Of course, any discussion of privacy in China would be incomplete without mention of the ways in which the Chinese government uses citizen data to implement its authoritarian goals.

Predictably, China's emerging consumer data privacy regime described above does not attempt to restrict the Chinese government's ability to surveil its citizens or access their data. All signs indicate that the government will continue to use intrusive techniques to surveil, monitor and oppress its

⁹ Sacks (2018), *China's Emerging Data Privacy System and GDPR*, <https://www.csis.org/analysis/chinas-emerging-data-privacy-system-and-gdpr>

¹⁰ See forthcoming analysis by Carrick Flynn (a CSET Research Fellow) for a more detailed treatment of the strategic role of data.

population, up to and including the techniques involved in the horrific treatment of Muslim Uighurs in Xinjiang.¹¹

These activities represent gross human rights violations, and deserve the attention of the U.S. Government.

However, a discussion of the future of AI is not the right venue for such attention. As much as the CCP would love to be perceived as having bleeding-edge AI surveillance tools at their fingertips, in reality the technologies that appear to be in use in Xinjiang and elsewhere (such as facial recognition or predictive policing) are straightforward applications of widely available data analysis tools. This is to say nothing of the even more basic methods at play, including ubiquitous checkpoints, the ability to seize and search cellphones, eavesdropping on electronic communications, and regular old human intelligence.

It would therefore be extremely challenging to effectively slow down Chinese access to these technologies. I fear that if the conversation about Xinjiang and other Chinese authoritarianism focuses too heavily on AI and other new technologies, the wrong countermeasures will be taken. While the impulse to ensure that U.S. researchers are not contributing to the surveillance state is a laudable one, we must not deceive ourselves that research like this is a key ingredient in China's actions. What's more, as described in the section on research collaborations above, we must be mindful of the costs to U.S. innovation and competitiveness that could come from poorly-targeted controls.

The determining factor in Chinese oppression is the CCP's willingness to pursue totalitarian ends, not the technological sophistication of its means.

As such, the goal of any measures to condemn the situation in Xinjiang should be just that—to condemn. Measures with the goal of preventing the development or use of a given technology will not work, and will instead damage the United States' standing on the world stage as the global leader in science and technology.

5. Standards and standardization

a. Chinese standardization efforts

China is well aware of the power that can be gained by having a hand in the design of widely implemented standards. As such, there has been an active push within China to develop and establish standards for AI.

One of the most prominent aspects of this push was the release of an in-depth white paper on AI standards in January 2018, which included contributions from over two dozen Chinese companies, associations, and academic organizations.¹² Another was a meeting held in Beijing in April 2018, where this white paper was presented to the first meeting of SC 42, a subcommittee on AI that sits within

¹¹ See, for example, this recent Human Rights Watch report for a detailed description of one strategy to collect and use citizen data in Xinjiang: <https://www.hrw.org/report/2019/05/01/chinas-algorithms-repression/reverse-engineering-xinjiang-police-mass-surveillance>

¹² Ding, Triolo, & Sacks (2018), *Chinese Interests Take a Big Seat at the AI Governance Table*, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/chinese-interests-take-big-seat-ai-governance-table/>

two internationally respected standards bodies, the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC).¹³

Due to the general purpose nature of AI technologies, it is likely to be more difficult and time-consuming to develop standards for these than for other technologies. So far, the efforts described above do not seem particularly close to producing specific technical standards that will be widely implemented, instead focusing on statements of broad ethical principles. Principles of this kind offer an opportunity for U.S. stakeholders to engage with their Chinese counterparts to reaffirm the importance of the ethical principles expressed, as well as providing a venue to point out incongruities between stated principles of this kind and how the technology is being utilized on the ground.

b. Informal standardization

It is worth noting in this context that while few formal, top-down standards yet exist for AI systems, the widespread use of specific platforms to develop and deploy AI models provides an analogous opportunity to influence the technology.

To the extent that such platforms exist, they primarily stem from U.S. companies. Prominent examples include two software libraries for deep learning: Tensorflow and Pytorch, developed by Google and Facebook respectively, are by far the most widely used platforms of their kind, including in China. (This despite attempts by Chinese companies to release their own versions, most notably Baidu's PaddlePaddle). Platforms like this provide the United States with a form of AI-relevant soft power, and in some ways could be considered analogous to a bottom-up standardization process, inasmuch as they could be used in certain circumstances to affect widespread features of the technology.

6. Recommendations

a. Specific policy recommendations

Specific measures Congress can take to strengthen U.S. competitiveness in AI and protect U.S. interests include the following:

- Improve immigration options available to AI researchers and engineers. As described above, foreign students studying in the United States attempt to stay at very high rates, but this could be threatened by strategic immigration policies currently being enacted by other countries in tandem with a worsening U.S. immigration environment. Specific options here include lifting numerical limits on H-1B visas and/or green cards for AI workers; creating a clear path from student/scholar status to permanent residence; and reducing processing times and application burdens. A forthcoming report from CSET will lay out immigration policy options to bolster U.S. competitiveness in AI in more detail.
- Allocate resources to the National Institute of Standards and Technology to support its efforts to develop and implement standards for AI, as initiated by the recent NIST Request for Information on this topic.¹⁴

¹³ *ibid*

¹⁴ Federal Register (2019), *Artificial Intelligence Standards*, <https://www.federalregister.gov/documents/2019/05/01/2019-08818/artificial-intelligence-standards>

- Increase R&D funding for basic AI research, for example by allocating new funding to the National Science Foundation's Directorate for Computer and Information Science and Engineering. Strength in fundamental research is the backbone of American advantage in AI, but no major federal effort has been made to strengthen that backbone during the current wave of progress in deep learning (in contrast to many other countries, especially China).

b. General recommendations


Beyond specific policy options, I also offer the following suggestions to the Commission and to Congress, to inform any other future action relating to AI competitiveness:

- Place liberal democratic values front and center, in action as well as in word. These values are what give the United States a sustainable, long-term advantage over countries like China. The more we compete on brute force, the better China's chances. If it is clear that we are competing on values, the whole free world is behind us.
- Relatedly, recognize the benefits the United States has drawn from being perceived by researchers as a good place to live and work, and seek to maintain or improve these attitudes. Do not lose sight of the wide-ranging benefits AI can bring to the United States and other countries alike, and have those benefits in mind when enacting measures that will affect how and where AI researchers can work.
- Where it is necessary to restrict foreign access to U.S. research, ensure that these restrictions are highly targeted and that their motivations are explained clearly. Work closely with experts with a strong understanding of the technology to ensure that the restrictions will not have unintended side-effects.
- Wherever possible, provide information rather than enacting restrictions. The AI research community is actively engaged in discussions about ethical and unethical uses of the technologies under development. The USCC, Congress, and other U.S. government entities can play an important role in informing that community about realities on the ground, for example about connections between Chinese research institutes and the Chinese state.

Once again, I wish to thank the Commission for this opportunity to speak on behalf of CSET and to address these issues with you. I look forward to your questions.

Figure 1¹⁵

Migration of Inventors, 2000–2010



¹⁵ Kerr, Kerr, Özden, & Parsons (2016), *Global Talent Flows*, <https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.30.4.83>