STANFORD-NEW AMERICA DIGICHINA PROJECT

AI POLICY AND CHINA
Realities of State-Led Development

Special Report No. 1

October 29, 2019

Edited by
Graham Webster
Foreword by
Andrew Grotto
and Ian Wallace

Stanford | Geopolitics, Technology, and Governance
Cyber Policy Center

NEW AMERICA

About DigiChina

The Stanford-New America DigiChina Project is a collaborative effort to understand China's digital policy developments, primarily through translating and analyzing Chinese-language sources. How the Chinese party-state and society deploy and use digital technologies is increasingly consequential for governance, markets, and security around the world. From laws and regulations to published commentaries from policymakers and corporate strategies, China's public sphere provides a great deal of insight into what's happening and why. DigiChina's network of contributors, working collaboratively with industry and policy experts around the globe, work to bring that insight to an international audience and inform important debates and decisions.

DigiChina has been supported through a partnership with the Ethics and Governance of AI Initiative of the MIT Media Lab and the Harvard Berkman Klein Center. In June 2019, DigiChina became a joint project of New America's Cybersecurity Initiative and Stanford University's Program on Geopolitics, Technology, and Governance.

About the Stanford Program on Geopolitics, Technology, and Governance

The Program on Geopolitics, Technology, and Governance (GTG) at the Stanford University Freeman Spogli Institute's Cyber Policy Center is dedicated to world-class scholarly and policy-oriented research on the political, legal, and economic implications of digital innovation and global competition. Artificial intelligence, quantum computing, and a proliferation of smart, connected devices will revolutionize warfare and create new challenges and opportunities in statecraft. They will enable automation in countless domains and lead to as yet unknown applications that catalyze new industries and business models. In the process, they will massively alter how economic value is created, captured, and distributed, with ripple effects in the domestic politics of nations and the broader global political economy.

These developments will impose profound demands on domestic and international political and regulatory institutions. Our research explores how the convergence of risk factors such as cybersecurity, privacy, and safety creates complex interdependencies across risk domains, and how governments and businesses craft policy and manage these risks in a global context to optimize the benefits of digital innovation.

About New America's Cybersecurity Initiative

New America is pioneering a new kind of think and action tank: a civic platform that connects a research institute, technology lab, solutions network, media hub, and public forum. New America's Cybersecurity Initiative aims to help create a cyber-secure America that lives up to its values at home and abroad. Our principal focus is to seek out opportunities to deliver impact at scale, particularly in areas that others have not been willing or able to address and that combine and focus the strengths of New America and our partners.

Our DigiChina partnership with Stanford University is one such project. More widely we believe that better cybersecurity will be achieved by a better flow of ideas and the engagement of a larger and more diverse group of people. We therefore aim to make New America a platform for good ideas and for information exchange— through our Fellows network, through events both public and private, and through our mentoring of interns and teaching of others.

Table of Contents

Foreword

A fun parlor game to play when traveling between Washington, D.C., and Silicon Valley is to ask people in both places to free associate words with China. In D.C., where the respondents trend toward politicos and policy wonks, the game generates a string of words along the lines of "competitor," "IP theft," "surveillance," and "military threat." The same game played with technologists and venture capitalists in Silicon Valley, however, yields words like "investor," "research partner," and "market."

The nature of the game primes respondents to offer simple answers, of course, but it does have a way of drawing out essential features of the debate about China in the United States and elsewhere, in the same way that caricature can offer incisive analysis about its subjects. In this case, the game highlights a cognitive dissonance between American policy and business leaders about China, while hinting at deeper layers of dissonance about the promise and risks of digital technologies and the role of government in optimizing this equation.

It is tempting to chalk this dissonance up to one great misunderstanding about China. And there is plenty about China and its ambitions, especially with respect to technological development, that is poorly understood—especially when it comes to the vast wilderness of digital policy challenges that Chinese policymakers are struggling to navigate, much like their counterparts in the United States, Europe, Japan and elsewhere. As one example, Siodhbhra Parkin's essay in this volume is a fascinating account of government and civil society initiatives in China to examine how AI might be deployed to better serve people with disabilities.

In debates in the United States and elsewhere about China's AI and related efforts, there is a tendency to treat China as a unitary, undifferentiated actor with crystal clear plans and an implementation glide path. The reality, however, is far more nuanced and complex, as demonstrated by DigiChina's past work and the essays in this volume from Lorand Laskai and Helen Toner, Mei Nelson, Jordan Schneider and Karman Lucero.

Reducing misunderstanding and misperception is a critical step toward more effective management of fundamental differences in values and policy orientations between China and the United States. One major difference that has become a sticking point in the trade conflict between the countries is China's multifaceted industrial policy. Essays by Benjamin Cedric Larsen on China's formal efforts to identify national champions for AI and Thomas Lehmann on a pair of local government initiatives to support AI research, development, and application each shed new light on different dimensions of Chinese industrial policy.

Johanna Costigan's contribution of interviews with a diverse group of four experts on China's AI and broader digital ambitions highlights the value and need for interdisciplinary work in China.

Graham Webster, who leads the DigiChina network of experts, has assembled an impressive volume that captures the core of what DigiChina is about: analysis of China's digital growth backed by real data in the form of careful translations of Chinese language materials. This report is also a direct result of the generous support of the Ethics and Governance of Artificial Intelligence Initiative of the Harvard Berkman Klein Center and the MIT Media Lab, which has cultivated a new crop of groundbreaking work on AI and society under the leadership of Tim Hwang.

This volume also marks the continued development of DigiChina as a collaboration between Stanford University and New America. There is an urgent need, now more than ever, for the contributions that DigiChina makes to research, teaching and policy work on China. We are excited to support Graham and the DigiChina network as they endeavor to take this work to new heights.

Andrew Grotto
Stanford University

Ian Wallace
New America

Introduction

BY GRAHAM WEBSTER

When China's government announced its ambitions for the country's theoretical, technological, and applied artificial intelligence development to reach a "world-leading level" by 2030, governments and markets worldwide took notice.

So did DigiChina. The New Generation Artificial Intelligence Development Plan (AIDP), drafted by experts across China's bureaucracy and issued by the State Council in July 2017, was one of this nascent project's first major translations. Our team of four translators then split up to provide three different views of its significance—as a legacy of central planning, "not a moonshot"; as a bureaucratic maneuver by its authors, but one with an "uncommonly foresighted approach" to AI governance challenges; and as a detailed plan that could portend "surpassing the United States."

Since 2017, Chinese officials, businesspeople, and researchers have mobilized remarkable efforts, even if the AIDP's authors might acknowledge their specific dozen-year targets were educated guesses.

In two years, especially in industry and in the market, significant advances in AI technology implementation have taken place. At the same time, real debate has come to the fore over how best to spur high-tech development. These Chinese efforts do not occur in isolation, however. Perhaps the biggest change since the summer of 2017 for Chinese tech development has come from the U.S. government. As U.S. sanctions enforcement activity has threatened first the telecommunications equipment company ZTE, then its much larger competitor Huawei, with being cut off from crucial U.S. components, the immediate continuation of China's development trajectory has been called into question.

This special report is the first of two volumes, with this collection assembling the insights of 14 researchers. This volume's first article, by Lorand Laskai and Helen Toner, documents how China's deep learning implementers, for the time being, depend on U.S. semiconductor technology potentially subject to export controls. China's national drive for "indigenous" or "independent" innovation is

in part intended to ease that sort of dependency, but the path ahead may be long.

The following four articles explore Chinese discourse around AI development efforts, revealing real debate, and describing two institutional maneuvers the government is executing in hopes of fueling development. Mei Nelson documents Chinese researchers debating the virtues of state-led innovation programs, versus bottom-up, market-driven innovation, revealing substantive disagreement about how to advance the nation's industries. Jordan Schneider takes the pulse of the lively online media space on AI, finding front-line businesspeople struggling with bad data and scarce—and picky—expert labor.

Benjamin Cedric Larsen identifies a potentially new form of AI industry governance as officials pick "National AI Team" companies for special privileges in given sectors, but deputize them to co-regulate their sector in aligntment withcompatible with party-state goals. Thomas Lehmann profiles two major AI initiatives in Beijing that could serve as a model for other localities seeking to build up a regional innovation cluster.

DigiChina emerged organically out of the reality of diverse viewpoints and interests in studying China's digital economy and technological development. We sought out four specialists who have engaged with AI and China in different ways and for different reasons, and Johanna Costigan interviewed them for our "From All Sides" feature.

Two final articles provide a jumping off point for Volume 2 of this report, which will document Chinese debates and efforts around AI ethics and governance. Siodhbhra Parkin identifies a Chinese governance gap that, properly handled, AI development could help fill: a holistic approach to serving and empowering people with disabilities. Finally, Karman Lucero observes that a rash of politically-charged central government ambitions for AI, combined with the Xi Jinping-era's broad political tightening, could result in an AI governance deficit as officials lack the freedom and confidence to experiment.

This report reflects DigiChina's mission of providing high-integrity, public spirited scholarship, to elevating emerging voices in the study of technology and society, and to assess Chinese events on their own terms. ◆

Graham Webster leads the Stanford-New America DigiChina Project at the Stanford University Cyber Policy Center's Program on Geopolitics, Technology, and Governance and New America's Cybersecurity Initiative, where he is also a China Digital Economy Fellow. He was previously a Senior Fellow and Lecturer at Yale Law School's Paul Tsai China Center, where he was responsible for U.S.–China Track 2 dialogues for five years. In the past, he wrote a CNET News blog on technology and society from Beijing and taught East Asian politics at NYU's Center for Global Affairs. Graham holds a master's in East Asian studies from Harvard University and a bachelor's in journalism from Northwestern University. He is based in Oakland, California.

Can China Grow Its Own AI Tech Base?

Despite market success, Chinese experts see stubborn dependencies

BY LORAND LASKAI AND HELEN TONER

Last December, China's top AI scientists gathered in Suzhou for the annual Wu Wenjun AI Science and Technology Award ceremony. They had every reason to expect a feel-good appreciation of China's accomplishments in AI. Yet the mood was decidedly downbeat.

"After talking about our advantages, everyone mainly wants to talk about the shortcomings of Chinese AI capabilities in the near-term—where are China's AI weaknesses," said Li Deyi, the president of the Chinese Association for Artificial Intelligence. The main cause for concern: China's lack of basic infrastructure for AI.

More than two years after the release of the New Generation Artificial Intelligence Development Plan (AIDP), China's top AI experts worry that Beijing's AI push will not live up to the hype. The concern is not just that China might be in for an "AI winter"—a cyclic downturn in AI funding and interest due to overly zealous expectations. It's also that for all China's strides in AI, from multi-billion dollar unicorns to a glitzy state plan, it still lacks a solid, independent base in the field's foundational technologies.

The concern seems counterintuitive at first glance. In recent years, China has built a crop of commercial AI juggernauts with no direct counterparts elsewhere in the world. Yet, upon closer scrutiny, it's clear that Chinese AI researchers are highly reliant on innovations and hardware built in the West.

Chinese Domestic Programming Frameworks Lag U.S. Giants

A brief glance at the infrastructure Chinese developers are using to run their algorithms reveals one reason for concern. The two dominant deep learning frameworks are TensorFlow and PyTorch, developed by Google and Facebook, respectively. A "framework" is essentially a set of programming shortcuts that makes it simpler for researchers and engineers to design, train, and experiment

with AI models. Most AI research and deployment uses one framework or another, because frameworks make it possible to use common deep learning concepts (such as certain types of hidden layers or activation functions) without directly implementing the relevant math.

While Chinese alternatives to TensorFlow and PyTorch exist, they have struggled to gain ground. Baidu's PaddlePaddle scarcely appears in either English- or Chinese-language listicles of top framework comparisons. Although it's difficult to find reliable and up-to-date usage statistics, various informal indicators all point to a large discrepancy in usage. According to Github activity, Baidu's PaddlePaddle trails PyTorch and TensorFlow by a factor of 3–10 on various statistics. In one Zhihu thread on comparing frameworks, only one user stood up for PaddlePaddle—the PaddlePaddle official account.

It's hard to imagine China's AI sector achieving the State Council's ambition to reach "world-leading levels" if the foundational software underlying its own research is built in the United States.

In the short term, the popularity of different frameworks may not matter much for AI research in China. But in the longer term, it's hard to imagine China's AI sector achieving the State Council's ambition to reach "world-leading levels" if the foundational software underlying its own research is built in the United States. What's more, the network effects that arise because researchers want to use the same frameworks as their collaborators (and because frameworks with more users are generally better maintained over time) mean it could be increasingly difficult for a Chinese company to come from behind and dethrone established frameworks.

No Clear Escape From GPUs or U.S.-Made Successors

When it comes to AI hardware, the outlook is equally troubling for China. Despite buzz in venture capital circles about Chinese AI chip startups like Cambricon and Horizon Robotics, Chinese AI developers continue to rely heavily on western hardware to train their neural networks. This is because Chinese AI chips have so far largely been confined to "inference," or running existing neural network models. In order to "train" those neural nets in the first place, researchers need high-performance, specialized hardware. Unlike most computational tasks, training a neural network requires massive numbers of calculations to be performed in parallel. To accomplish this, AI researchers around the world rely heavily on graphics processing units (GPUs) that are mainly produced by U.S. semiconductor company Nvidia.

Originally designed for computer graphics, the parallel structure of GPUs has made them convenient platforms for training neural networks. SenseTime's supercomputing center DeepLink, for instance, is built on a staggering 14,000 GPUs. However, GPUs are not the only hardware platform that can train neural nets. Several chips including Google's Tensor Processing Unit (TPU) and field-

programmable gate arrays (FPGAs) from companies like Intel and Xilinx will likely reduce the importance of Nvidia GPUs over time. Notably, none of these competitors to the GPU are Chinese.

Why are there no Chinese competitors challenging the GPU's reign? The answer, according to Sun Yongjie, a notable tech blogger in China, is that Chinese AI chips are created for "secondary development or optimization" rather than replicating fundamental innovations. The derivative nature of Chinese AI startups came into stark relief last year when California-based Xilinx bought DeePhi Tech, a trailblazing Chinese AI chip startup. The acquisition provoked immediate indignation among netizens, many of whom argued that the Chinese government should have intervened to protect one of China's most promising chip ventures. Upon further reflection, however, several bloggers argued government intervention would be fruitless, since DeePhi's deep learning processors are entirely built on Xilinx FPGA frameworks. "If DeePhi Tech ever broke away from Xilinx's FPGA platform, it would be completely cut off from all sustenance," Sun wrote.

DeePhi's technical dependence on a western chip company is not an anomaly—it's the industry norm. Horizon Robotics, China's largest AI chip unicorn, often billed as the "Intel of China," built its main AI processor architecture, the Brain Processing Unit (BPU), on top of Intel's FPGA. Most Chinese AI companies buy the license for core components, rather than developing them internally. In some cases, according to industry insiders, AI startups have even outsourced the actual chip design to more experienced western design companies.

Breakthrough vs. Implementation

It's an open question whether any of this matters. Does China need to develop its own foundational software and hardware in order to be an AI leader, or can it build upon the existing scaffolding of western companies? If Chinese AI researchers can effectively use TensorFlow and train models on Nvidia GPUs, does it matter whether foundational platforms are also built in China?

At least one prominent voice—venture capitalist and AI scientist Kai-Fu Lee—thinks it does not. In his recent book *AI Superpowers: China, Silicon Valley, and the New World Order*, Lee argues that AI has entered the "age of implementation." The foundational breakthrough of modern AI research (deep learning) has already been made, Lee claims, so now all that matters is translating and applying that breakthrough for specific use cases. In many ways, China's approach to AI thus far appears to be based on this premise, with companies flexing their implementation muscles to reach giant valuations and fast adoption.

But Lee's view is far from consensus among AI researchers.

The basic algorithmic ideas behind modern deep learning systems were in place by the 1980s, but constraints on data and computation made them impractical to implement until more recently. Experiments in 2011 and 2012 kickstarted the current deep learning boom, using the method to achieve state-of-the-art results in image recognition. Like most interesting advances in AI research, the novel contributions of these experiments were neither major breakthroughs nor minor tweaks of past research—they were something in-between.

Looking forward, there are plenty of reasons to expect a steady stream of medium-sized advances to continue pushing the bounds of AI. Areas that have seen significant progress in the last year or two include image generation, increasingly complex strategy games, and—most recently—language understanding and generation. The research powering each of these advances came out of labs focused on fundamental R&D, not mere "implementation."

A consensus in the Chinese AI community has gradually formed around the view that China needs to participate in this steady foundational progress in order to become an AI powerhouse. The catalyst for this convergence in thinking was largely external: The April 2018 addition of Chinese telecommunications company ZTE to the U.S. Commerce Department's Entity List, which lists companies and other entities U.S. firms may not export certain items to without a special license. In a speech before the National People's Congress Standing Committee that fall, Chinese Academy of Sciences (CAS) expert Tan Tieniu* warned that China's AI industry could face its own ZTE moment if it did not build its own foundational technology. Far from the isolated opinion of a risk-averse government scholar, the fear resonated throughout much of China's AI community. A post on DeepTech, a popular WeChat account for AI industry news, called TensorFlow and other U.S. open-source frameworks "traps" that could "suffocate" China's AI development, asking: "If open-source projects can be curtailed on a whim by export bans, will China's AI companies be next?"

Aside from the threat of U.S. export control, prominent members of China's AI community have become more vocal about the constraints that a lack of independent foundations could place on China's AI development. "As the AI era progresses, the constraints on our AI industry from AI algorithms and computational power, especially AI chips, have become clear," noted AI expert and Baidu executive Wang Haifeng said.

Starting from Scratch

What would it require for China to patch its foundation deficit? Naturally, reinventing the wheel—or in this case reinventing TensorFlow—will not do. Rather, Chinese researchers will need to push the bounds of basic AI research, contributing new ideas to the global research community and building their own foundational platforms. The problem is that China, despite its strides in commercializing AI, does not appear to be making much progress in basic

*Read more from Tan's speech on DigiChina:

"…From a technical perspective, AI is currently at a technical inflection point: going from 'unusable' to 'useable,' but the distance to 'useful' still has bottlenecks such as those related to data, energy consumption, generalization, interpretability, reliability, security, etc. …'

(Translated by Cameron Hickert and Jeffrey Ding)

research. According to data compiled by Elsevier in partnership with CAS, the "citation impact" of Chinese AI papers remains significantly behind the numbers from Europe and the United States. Anecdotally, deep learning researchers point to one significant contribution from a Chinese lab in recent years—a paper out of Microsoft Research Asia's Beijing lab in 2015 introducing "residual networks," a training technique now widely used by other researchers —but most people struggle to name a second.

The most obvious reason for China's struggles with basic AI research is old news: brain drain. Although efforts to train more Chinese AI researchers have succeeded, a huge fraction of those researchers—by one estimate, almost three-quarters—end up overseas. This stands in stark contrast to the United States, which is a massive net importer of AI talent. A major reason for this imbalance is the high quality of research labs in the United States, in both academia and industry. U.S. companies like Google and Facebook and universities like MIT and Stanford regularly top the charts of labs producing the highest volume of papers accepted at top conferences. (Perhaps not coincidentally, the lead author of the residual networks paper mentioned above has since moved from Beijing to a Facebook lab in California.)

China's tech sector still has no credible competitor to Windows operating systems or the Intel CPUs. The question is, will China's push to patch its foundation deficit in AI be different?

Building up cutting-edge research capacity is a chicken-and-egg problem: The best researchers want to work with other outstanding researchers, giving a natural advantage to established labs and making it hard to bootstrap a great lab from scratch. High salaries and other incentives are often insufficient to overcome this dynamic, as demonstrated by Baidu's struggles to retain high-profile AI talent.

Commercial incentives in China's AI industry also reinforce the country's basic research deficiency. Chinese tech companies largely underinvest in R&D, pursuing commercial applications over basic research. While the hype of AI has triggered a surge in attention from investors, investments have been mainly focused on applications that can be quickly commercialized. According to a survey conducted by EO Intelligence, a Chinese market research firm, AI ventures working in financial technology, service industries, and surveillance have received a disproportionate amount of investment. At the same time, ventures working on foundational components have been largely ignored by investors.

These incentives, according to voices in the Chinese AI community, have perversely influenced what AI researchers in China work on. "There are a lot of feasible projects, but most do not generate excitement. Experts pick the hottest project—this reflects a type of low self-esteem," said Lu Ruqian, a CAS scholar and early pioneer in AI. "I believe this blind herd mentality is creating a

dangerous situation," said Han Liqun, a professor at Beijing Technology and Business University.

Indigenous Innovation 2.0?

AI is not the first area of technology where Chinese experts have identified a national deficiency. Amidst the widespread adoption of information technology in the 1990s, government scholars bemoaned China's reliance on foundational western IT systems and called for a concerted effort to build homegrown alternatives and wean China off western technology. This effort to harness "indigenous innovation" was placed front and center in the 2006–2020 National Medium- and Long-Term Plan for Science and Technology Development. Decades later, the results have been lackluster: China's tech sector still has no credible competitor to Windows operating systems or Intel CPUs. The question is, will China's push to patch its foundation deficit in AI be different?

One major difference lies in the commercial sector. The ZTE incident appears to have aligned industry with the government on the importance of building foundational technology. The addition of Huawei to the entity list this past spring might have been a second catalyzing event, once again reinforcing this point.

"If we do not master the core technologies, we will be building roofs on other people's walls and planting vegetables in other people's yards," said Alibaba founder Jack Ma shortly after Huawei was targeted. In September 2019 Alibaba's chip subsidiary Pingtouge released its first dedicated AI processor for cloud computing. A month earlier, Huawei announced its first AI training chip and first open-source deep learning platform, MindSpore.

These developments suggest that the critical variable in whether China can catch up might not lie within China but in the United States. If the United States closes off its AI ecosystem, Chinese AI researchers have a meaningful incentive to develop their own platforms, no matter the costs. ◆

Lorand Laskai is a J.D. candidate at Yale Law School and a Visiting Researcher at Georgetown's Center for Security and Emerging Technology (CSET). Previously, he worked as a Research Associate at the Council on Foreign Relations and a Lead Researcher at the Financial Times in Beijing. His writing on China and its approach toward cyberspace and high-tech development has appeared in publications like *Foreign Affairs*, Slate, China Brief, and *Foreign Policy*. He holds a B.A. from Swarthmore College.

Helen Toner is Director of Strategy at Georgetown's Center for Security and Emerging Technology (CSET). She previously worked as a Senior Research Analyst at the Open Philanthropy Project, where she advised policymakers and grantmakers on AI policy and strategy. Between working at Open Philanthropy and joining CSET, Helen lived in Beijing for nine months, studying the Chinese AI ecosystem as a Research Affiliate of Oxford University's Center for the Governance of AI.

Shall We Take an 'Alternative Route,' or 'Overtake on the Curve'?

Debating the state-led development model for AI in China

BY MEI NELSON

Chinese strategists in AI fields are well aware of their country's limitations. In May 2019, at the China International Big Data Industry Expo, Chinese Academy of Engineering Academician Gao Wen gave a keynote speech identifying four advantages and four shortcomings in China's AI development. Among China's advantages, Gao said, were strong policy support, ample data, rich application scenarios, and many young people with potential talent. Gao identified shortcomings in: basic theoretical research and original algorithms, core AI components, open source platforms, and high-end talent.

Chinese planners had already summarized these shortcomings in the State Council's 2017 New Generation Artificial Intelligence Development Plan (AIDP). Gao's address suggests that China has made little significant improvement in the two short years since then when it comes to areas such as basic research, algorithms, or core AI components. Having acknowledged these weaknesses, the question before the Chinese AI community now is to address them.

China's AI development, like its other state-led science and technology development initiatives, attempts to combine "innovation-driven development strategy" with "top-level design." The Chinese government has led the drive to lead the global AI sector by 2030, under a changing set of campaigns, from "intelligent manufacturing," first mentioned in the Made In China 2025 strategy in May 2015, to the Internet Plus strategy of July 2015, which listed AI as one of 11 key action areas, to the 2017 AIDP.

While cheering the progress that Chinese AI efforts have made in such a short time, some Chinese scholars and industry experts have recently debated whether a state-run system is best for AI development and what alternatives could be better. Two recent examples illustrate the lively dialogue in China about how best to foster development of AI.

Top-level Design vs. State-run System

China's state-run system has created incentives that hinder innovation, argued scholars in the China AI Index 2018 report, jointly published in March 2019 by the Center for AI and Institutions (CAII), part of the prominent Cheung Kong Graduate School of Business, and the Big Data and Cloud Computing Lab of Wuhan University.

At the report release conference, experts pointed out that Chinese researchers published at least as many AI research papers as did U.S. researchers between 2008 and 2012, but questions have lingered about the quality and influence of those publications. Numerous papers by Chinese researchers had zero citations, while the number of papers with over 1,000 citations by Chinese-only authors is also much smaller than that of papers co-published by Chinese researchers together with U.S. researchers. Xu Chenggang, the director of CAII, blamed the incentive mechanism within the Chinese academic system: Many researchers are evaluated by the quantity of papers they publish, not the quality.

Resources for basic research should be determined by experts in the industry, and resource allocation should be determined by venture capital. The government should step back and let the market play its role, Xu said.

Going deeper, Xu identified the state-run system in AI development as the root cause, especially in basic research and the development of core AI components such as high-performance integrated circuits (or chips) for machine learning. A state-run system does not work in AI development, Xu argued. In a state-run system, the government decides how to allocate resources in all aspects. However, in a new and fast-developing field with high uncertainty such as AI, Xu argued, the resources for basic research should be determined by experts in the industry, and resource allocation should be determined by venture capital. The government should step back and let the market play its role, Xu said.

At the same conference, Liu Yadong, editor-in-chief of the *Science and Technology Daily*, also commented that the role of the government for successful AI development should not be leadership, but guidance. Specifically, Liu said, the government should not focus on planning, but on setting up rules within which experts and private companies can compete fairly.

Overtaking on Curves vs. Alternative Routes

Another area of ongoing debate is how China can best advance in the ranks of global AI developers. As stated in several national strategies, "indigenous" or "independent" innovation should be the driving force to improve basic research and develop AI-related core technologies, implying an innovation path less dependent on existing or future foreign technology. The emphasis on indigenous innovation first appeared in 2006-2020 National Medium- and Long-Term Plan

for Science and Technology Development (MLP) and later in the National Innovation-Driven Development Strategy. For AI development particularly, the 2017 AIDP further emphasized the goal of pursuing indigenous innovation in AI technology.

Several visions for how to achieve indigenous innovation have emerged. More than a decade ago, the MLP explained how indigenous innovation combines original R&D with "integrated innovation" and "**re-innovation**"*—by absorbing, reassembling, and upgrading imported technologies in new ways to produce original breakthroughs.

Since the MLP came out in 2006, however, China's approach to innovation has evolved. China's development of high-performance chips, for example, exemplifies indigenous innovation in practice, but also attempts more original innovation through a bottom-up approach and finding new ways to make breakthroughs. The common metaphor for this bottom-up approach is "**overtaking via alternative routes**."* Chinese technology experts such as Alibaba's Jack Ma applaud "overtaking via alternative routes" as a road to true success.

Sun Ninghui, director of the Institute of Computing Technology (ICT) at the Chinese Academy of Sciences (CAS), discussing China's chip development in a 2016 interview, explained three parallel but complementary development strategies, all promoted by the government and industry experts: re-innovation based on assimilation, absorption, and upgrade of imported technologies; innovation through overtaking on curves; and innovation through overtaking via alternative routes.

To the familiar "re-innovation" and "alternative routes" tactics, Sun added "**overtaking on curves**,"* which refers to catching up with foreign advances and looking for opportunities to pull ahead along the same technological track. This contrasts with an overtaking via alternative routes approach, in which one does not compete with adversaries in the same track but develops one's own independent technology. To illustrate, Sun gave the example of Cambricon Technologies, a start-up manufacturing AI processors with investment from the CAS's ICT and the e-commerce giant Alibaba. In 2016, Cambricon released China's first AI Chip, Cambricon-1A, which it claimed was the world's first commercialized neural network processor chip. Cambricon chips appear in smartphone chipsets produced by major Chinese manufacturers. "Cambricon avoided competing with Qualcomm and Microsoft in the current market, but was able to take the lead in AI field," Sun said. They took an "alternative route."

Because companies like Cambricon that can make breakthroughs in high-performance chip R&D field are very few, however, the other two development strategies are also viewed as necessary. Whereas "alternative route" innovation is

<div style="border: 1px solid #0a3d62; background:#0a3d62; color:white;">

*

Pick your innovation path:

- "Introduce, digest, absorb, re-innovate" (引进，消化，吸收，再创新): Assimilate what's out there, master it, and top it.
- "Overtake by alternative routes" (变道超车): Build it your own way, and beat the competition.
- "Overtake on the curve" (弯道超车): Learn the competition's game, but find a way to take the lead.

</div>

11

bottom-up, both "overtaking on curves" and "re-innovation" are top-down, state-backed incremental innovation models.

The development of the Loongson CPU chip, which based its design on the non-mainstream MIPS architecture by a U.S. fabless semiconductor design company (i.e. one that designs but does not build chips), exemplifies "overtaking on curves," whereas Sugon's joint venture deal with U.S. chip developer AMD exemplifies the "re-innovation" approach.

While many AI industry leaders push to innovate through the bottom-up strategy of overtaking via alternative routes, China's core of indigenous innovation remains centered around an advanced imitation strategy officially summarized as "introduce, digest, absorb, re-innovate." This old thinking, reflecting definitions from the 2006 MLP, was evident in one of the requirements for the prestigious AI science and technology award named for AI pioneer Wu Wenjun. The Enterprise Technology Engineering Innovation Program Award in 2018 called on applicants to "master core technologies and independent intellectual property rights and solve the key, common, and supporting technology problem for industry development through original innovation, integrated innovation, and introducing, digesting, absorbing, and re-innovation."

These debates over the best approach to advancement in AI-related industries reveal considerable dynamism below the surface of top-level goals and the consensus around which challenges China faces. Incentives may be mixed. If government awards value "re-innovation" while some experts favor "alternative route" development, a researcher could reasonably believe a bottom-up approach is most likely to work but still opt for the top-down path to be politically safe. Even more difficult, a shift to market-driven resource allocation would entail a break-up of the state-run system that is highly unlikely under current political conditions. ◆

Mei Nelson is an independent researcher. Her research focuses on East Asia cyber threat analysis and political, strategic, and economic affairs supporting various U.S. government organizations and commercial customers. Mei is a board member of RSA Conference Security Scholar program. She holds an M.S. in Management of Secure Information Systems from George Mason University.

An 'AI Winter' on the Horizon?

In the trenches, sometimes it's hard to live up to the hype

BY JORDAN SCHNEIDER

Is an "AI Winter" coming to China? Some leading investors in the Chinese AI space think so. This article looks at Chinese sources from two popular new media outlets to explore why some see stagnation where others see a booming industry.

DougLong and the Doldroms of the AI business

"DougLong," an anonymous investor and entrepreneur in AI, reflecting on Gary Marcus' recent book *Rebooting AI*, points to three big issues with AI in 2019. First, good data is hard to come by. Second, test training data doesn't line up with actual operating environments. And third, "to B" companies having a hard time retaining AI talent for the long haul.

While there is little available as to the DougLong's true identity, the piece was published by respected tech media outlet and consulting firm Jazzyear, giving it a large audience in a country where pseudonymous commentary is a common tool of discourse. On the first issue, the dearth of good data, DougLong quotes an anonymous old hand in the data industry who just joined an AI startup and has been sorely disappointed with AI's effectiveness in the marketplace.

> "When customers see how many terabytes or petabytes they have in their database, they think they have big data. … But once we get down to work, the data is basically useless. Some fields are mis-entered, others are too sparse. Once you finish cleaning it up, the data leads to totally logically unsound conclusions, with no chance to do any deep learning. …
>
> For instance, once I did a fault detection project for a Zhejiang tire factory. … Hundreds of thousands of tires were piled up in open air collecting dust, so we had to hire people climb these tire mountains, clear away the dust, and write down the tires' model and batch

numbers and their faults. But on hot days—dirty and tired—some workers just lazed around and wrote up fake data. …

The fact that some data sources are borderline illegal is an open industry secret. In some industries where information security measures are weak, it's most cost-effective to find some internal personnel to just copy off full hard drives."

DougLong used an analogy to Chinese medicine to explain the importance of AI engineers sticking to one project for an extended period of time. "Since the theory around deep learning isn't complete and it's impossible to understand the operating mechanisms of algorithms, the success of various adjustments and improvements depends on experience combined with luck, and capabilities are hard to replicate quickly. It's like learning Chinese medicine. For a novice practitioner to mature into a high-level talent, one must complete many cases and encounter many conditions. Old Chinese medicine practitioners accumulate personal experience of successes and failures, depending on cases and insights to accumulate experience in the '[four methods of diagnosis](#).'"

'We had to hire people climb these tire mountains, clear away the dust, and write down … their faults. But on hot days—dirty and tired—some workers just lazed around and wrote up fake data.'

But the way the industry is set up makes it difficult to support top talents, especially in "to B" businesses. DougLong writes that most projects require specialized, non-scalable work, and many senior people already have families and don't like long business trips. Said one cloud sales rep at a BAT (Baidu, Alibaba, or Tencent) firm:

"One time a customer asked us to do a proof-of-concept for an AI project and wanted some high-level people on it. So I pulled out all the stops and borrowed a few people from AI research institutions. They went on-site for six weeks, but the project didn't succeed. When I tried to get them again, they wouldn't answer.

They don't like doing client projects, and what's more, they can't use that time to publish papers. And with such expensive outlays of human resources, there's no guarantee of standing out in the year-end results."

Building Foundations, or Seeking Handouts?

This past August, a handful of senior investors reflected on the challenges facing domestic AI startups at a [conference](#) organized by the digital media outlet [Lieyun Net](#).

Liu Shui, investment director at the incubator [CAS Star](#), encouraged startups to focus on the tech underlying advances in AI. According to Liu, "AI chips are the main battlefield, regardless of whether you're pursuing hardware or algorithms." But he sympathized with the plight of entrepreneurs in the space. "Even if your

technology is top notch, it's very difficult to commercialize a product. First, finding funding is hard. And second, there isn't a lot of support to help transitioning from tech to a viable product." He also admitted that China lacks investors with a background in AI to help back the best startups.

Wang Sen'ao, an executive at Lieyun Net, advised AI startups struggling to find cash to "change yourself to let the government back you unconditionally." It's possible for a startup after a year of operation to get subsidies approaching U.S. $1 million, but doing so is different than pitching yourself to private sector investors. According to Wang, firms need a model that aligns both with government aims and their own interests. Doing so requires entrepreneurs to "push [government] guiding documents from top to bottom," because, Wang says, when it comes to national industrial policy, State Council instructions filter all the way down, and their administrative measures and implementing rules require rapid study.

Overall, independent Chinese AI firms are facing similar headwinds to western ones as they struggle to commercialize their technology. It remains to be seen whether government support is a difference-maker. Ultimately, DougLong leaves readers with the following advice: "Give up fantasies, buckle down, and stay strong to get through the winter." ◆

Jordan Schneider is the Beijing-based host of the ChinaEconTalk **podcast**. He also writes the ChinaEconTalk **newsletter** which puts out weekly translations of tech-related trending WeChat articles. He is on twitter at **@jordanschnyc** and his Chinese landscape paintings "show promise."

China's National AI Team

The role of National AI Open Innovation Platforms

BY BENJAMIN CEDRIC LARSEN

China's government has become increasingly active in supporting a national agenda of AI development. In doing so, it is devising new means of guiding development in greater concert with leading private sector enterprises that advance key AI technologies and applications. Mimicking previous development strategies, a few companies have been selected as "National AI Team" members, an endorsement that carries both national and local government support, as well as access to regional projects and related public data resources. In return, the government expects that key standards for AI ecosystem development can be coordinated with greater efficiency among stakeholders, while smaller enterprises are enabled to synchronize with leading AI developments through open innovation platforms.

The notion of "National New Generation Artificial Intelligence Open Innovation Platforms" (AIOIPs) originated in November 2017, when China's Ministry of Science and Technology (MOST) endorsed four private sector companies to construct platforms for specific purposes. The companies were Baidu (for autonomous driving), Alibaba (smart city), Tencent (medical imaging), and iFlyTek (smart audio, i.e. natural language processing). A fifth AIOIP, SenseTime (smart vision), was added in 2018.

In August 2019 the initiative was expanded to include 15 AIOIPs, and it remains open for further entities to apply. Applicants are assessed by a team of experts organized by MOST, while selection criteria rests on the applicant's technological capabilities and anticipated results of the AIOIP. Each applying entity has to pre-specify a subdomain of AI platform development, which the entity will focus on opening up to a broader array of companies for further interaction and development. While each sub-domain is expected to relate to a distinct area of AI development, it is clear that conceptual overlaps exist between the domain boundaries.

The AIOIP initiative relies on leading enterprises to promote deep integration of AI with the real economy, while companies are expected to deliver on four key tasks across: R&D, ecosystem participation, sharing data and open-source software, and supporting the entrepreneurship of small and medium-sized enterprises (SMEs). MOST, along with provincial-level S&T management departments, actively support the AIOIP initiative, for example by promoting

Autonomous Driving	Smart City	Medical Imaging	Smart Audio	Smart Vision
Baidu 百度	Alibaba Group 阿里巴巴集团	Tencent 腾讯	科大讯飞 iFLYTEK	商汤 sensetime
Vision Computing	**Smart Marketing**	**Software / Hardware**	**Inclusive Finance**	**Video Perception**
依图 YITU	明略科技 MININGLAMP Technology	HUAWEI	中国平安 PING AN 保险·银行·投资	HIKVISION
Smart Supply Chain	**Image Perception**	**Cybersecurity**	**Smart Education**	**Smart Home**
JD 京东	MEGVII 旷视	360金融 理财 贷款 保险	TAL 好未来	mi 小米 xiaomi.com

application of the relevant technologies. This implies providing space for testing autonomous vehicles (Baidu), access to city infrastructure for monitoring and upgrade (Alibaba), access to public medical data (Tencent), access to the judicial system (iFlyTek), and access to surveillance systems (SenseTime), among other public areas that are being opened to private-sector development. In return for official support, each AIOIP entity is expected to provide annual reports summarizing the ongoing progress of their open innovation platforms.

Opening Public Data While Supporting the Platform Economy

Through these collaborations, a new model of AI development and an associated governance model are emerging in China, where government-designated platforms and related public-private partnerships emphasize an experimental, gradual, and decentralized approach to selectively opening public domains and associated data repositories.

In the process, leading private sector enterprises are endorsed to apply innovative AI solutions to optimize public institutions and the provision of public goods and services, often implemented on a local and regional basis. Tencent's Miying platform, for example, partners with Shenzhen Hospital Center, where residents can undergo AI-based remote screening for retina problems related to diabetes, aiming to alleviate strains on the public healthcare system.

In August 2019, the State Council issued a guiding opinion on Promoting the Standardized and Healthy Development of the Platform Economy, which also calls for greater data sharing between government departments and platform companies. The guiding opinion views the platform economy as a new organizational mode of productivity and economic development, while regulatory oversight from the government is to be devised in greater concert with leading platform operators. The platform economy document and the

AIOIP initiative should therefore be viewed as congruent, emphasizing a larger structural role to be played by leading AI enterprises and other players operating in the platform economy.

Platform Leaders Devise Standards and Create the Rules for Ecosystem Engagement

The 15 enterprises that have been selected to establish open innovation platforms are, to a certain extent, crafting the structural rules that affect wider industry engagement. As AIOIPs enable start-ups and SMEs to enter and participate in ecosystem development, they do so through open access to data, toolkits, libraries, frameworks, computing resources, and sometimes competitions, which are accessible through Application Programming Interfaces (APIs) on the open innovation platforms. This suggests that the AIOIP initiative is less about granting preferred access to a few select companies, and more about enabling structural mechanisms that afford greater participation and innovation in emerging ecosystems and sectors that increasingly will be powered by AI technologies.

For Alibaba's ET City Brain, which cooperates with several local governments to provide smart city solutions from transportation to energy, water supply, and so on, this implies that Alibaba is responsible for ensuring that ecosystem partners and developers stay compliant with existing legal frameworks on areas of critical public infrastructure and information systems.

Companies and regulators are jointly charged with creating new standards for data pooling and interoperability between different but increasingly interconnected public and private systems.

As the selected AIOIP companies are all leaders in their respective fields, they are also the de facto architects of system-wide standards and interfaces, which often are shaped in collaboration with research institutes, universities, and policymakers. In the process, companies and regulators are jointly charged creating new standards for data pooling and interoperability between different but increasingly interconnected public and private systems.

iFlytek, for instance, cooperates with the Shanghai courts to develop Project 206, which seeks to upgrade case-handling in the judicial system. iFlytek also collaborates with educational institutions across 10 provinces to provide voice recognition technology for high school oral examination assessments. Several of China's National AI Team members also provide educational materials and establish AI curricula for use in China's educational system.

Building China's Open Source Resources

In terms of data and software sharing, several of the nationally endorsed AI platforms are also behind a culture change toward open source, echoed in the

release of deep learning frameworks including Baidu's PaddlePaddle, Alibaba's XDL, SenseTime's SenseParrots, and Huawei's MindSpore.

Companies, as well as planners, hope that opening AI frameworks to wider use through AIOIPs will unlock greater network effects, benefitting independent developers and SMEs, as well as expanding entrepreneurial activity in the wider ecosystems.

The release of proprietary frameworks by leading AI enterprises, not only marks a new beginning for China's open-source community but also may decrease China's reliance on existing frameworks and operating systems, which predominantly originate in the United States and may become less dependable if technological "decoupling" deepens. Huawei's commitment to rapidly develop its HarmonyOS mobile operating system, as well as the long-standing pivot towards engineering indigenous AI chipsets, should also be viewed in this regard.

Pursuing National Goals Beyond State-Owned Enterprises

The AIOIP initiative extends an opportunity for private firms to rethink the link between public and private sectors. Traditionally, partnerships involving strategic assets and sensitive information have been reserved for China's state-owned enterprises, but policymakers have realized that rapidly advancing AI capabilities are developed outside the immediate scope of traditional SOEs. This new approach to AI development aims to move in lockstep with the advancements of leading AI companies and platforms, which emerge as regulatory stakeholders in the process. In doing so, the capabilities of private firms are nudged closer to the Chinese government's long-term visions, such as those articulated in the New Generation AI Development Plan (AIDP).

Parts of the initiative, however, are indicative of a broader international trend in which governments increasingly provide APIs to trusted intermediaries as a way to open public infrastructure and data repositories to private sector development. Framed as Government-as-a-Platform (GaaP), the public sector can act as a catalyzer for AI development through new procurement strategies, in which the public sector is able to incentivize and subsidize the development and innovation of AI systems for the government's vision of the public good.

In this new approach to AI development, leading companies and platforms, emerge as regulatory stakeholders.

Meanwhile, by helping establish open standards and interoperability between data and AI systems in areas where the government is opening up, policymakers can forestall the risk of single-player market dominance by encouraging broad systemic operability. In China, these tendencies are currently seen, as many of the enterprises behind AIOIPs compete across each other's designated domains.

From a business perspective, the establishment of AIOIPs mirrors developments of leading AI platforms elsewhere, and would likely have taken place regardless of central planning. In terms of policy, however, the AIOIP initiative embodies a new approach to the development of AI, which is symptomatic of China's long-term strategic thinking on the role of AI in the broader economy.

In the United States, the Computing Community Consortium's 20-Rear Roadmap for Artificial Intelligence, released August 2019, perhaps comes closest to the AIOIP vision. The roadmap calls for sustained support from the U.S. federal government and prescribes that the government should focus on establishing an open AI platform, to open resources such as data and machine learning libraries, and to establish national AI competitions, national AI research centers, and AI curricula for the education system.

While China already is implementing its own approach to AI development, the intricate relationship between private sector enterprises and policymakers is likely to come under increased scrutiny as its economic and technological dispute with the United States continues. For China's National AI Team, the political drive to establish AIOIPs may prove to have adverse consequences for these companies expansion in advanced markets. Nonetheless, the AIOIP initiative exemplifies an innovative approach to supporting the development of AI in a budding platform economy, while it remains to be seen just how effective the approach will be, in terms of spurring wider ecosystem participation and AI innovation. ◆

Benjamin Cedric Larsen is a Ph.D. Fellow at Copenhagen Business School (CBS) and Chinese Academy of Sciences (CAS), as well as a Visiting Researcher at Stanford's Asia Pacific Research Center. Benjamin holds a dual master's degree in public management from CBS and CAS in Beijing. Prior to pursuing his doctoral degree, Benjamin worked professionally in China, holding positions across the Danish Foreign Ministry as well as China's technology industry (NJU Electronics / JD AI). Benjamin's research interests relate to the organization of AI innovation, with a particular focus on the dynamics playing out in China. His research looks at the organization of leading AI ecosystems while highlighting the role of the state in the governance of digital platforms that shape innovation in AI.

AI Politics is Local

Municipal and central authorities
Dream up a model AI cluster in Beijing

BY THOMAS LEHMANN

China's central government plans for national advances in AI industries call for widespread efforts across numerous industries and in academia. The 2017 New Generation AI Development Plan (AIDP) further calls for encouraging "local industry chains and innovation chains around AI" and building "AI industry clusters." Local institutions across the country have announced such efforts, and the Beijing area, already home to prominent companies and research institutions, is perhaps a prime example.

In November 2018, China's Ministry of Science and Technology and the Beijing Municipal Government announced the Beijing Zhiyuan Action Plan (the "Zhiyuan Plan") and established the Beijing Academy of Artificial Intelligence (BAAI). BAAI was formed from a coalition of academic and private sector leaders, with backing from some of Beijing's most influential institutions and corporations in AI, including Peking University, Tsinghua University, the Chinese Academy of Sciences, Baidu, ByteDance, Megvii, Meituan-Dianping, and Xiaomi.

Here, we provide profiles of Zhiyuan Plan and BAAI, two local initiatives that stand out from the pack because of central government support and location in one of China's best technologically endowed cities.

Beijing Zhiyuan Action Plan

The Zhiyuan Plan is a response to a comprehensive web of central and municipal science and technology development policies and initiatives. Where the AIDP provided broad guidelines and goals for the AI industry, the Zhiyuan Plan makes the central government's directives more concrete.

According to Beijing Municipal Science and Technology Commissioner Xu Qiang, the Zhiyuan Plan is meant to carry out the following four responsibilities:

1. Create an innovative AI ecosystem, build Beijing's open source platform, and promote open source algorithms through the use of public data, smart computer programming frameworks, and computing infrastructure;
2. Construct a high-level joint lab to address core basic ethics questions, launch integrated and collaborative research, and promote indigenous innovation;
3. Identify, convene, and cultivate elite AI talent; and
4. Establish Beijing as a global hub through strengthening corporate, academic, and institutional cooperation, as well as holding global AI summits.

No major publicly released document is attached to the Zhiyuan Plan, but correlated projects have included the establishment of BAAI and initiation of the Beijing Zhiyuan Scholars Plan, the National New-Generation AI Innovation & Development Pilot Zone (the "Pilot Zone"), and the Beijing Zhiyuan-Megvii Intelligent Model Design and Image Perception Joint Lab (the "Joint Lab"). The Pilot Zone is the first of its kind in China, while the Joint Lab is part of an evolving institutional domestic model. Over two years, an increasing number of corporate and academic joint AI labs have been established around China, as well as between China and other countries.

Despite this activity, measurements of success are not yet available. The AIDP set out timelines and goals around core AI and AI-related industry scales over the next 15 years. Other measurements of success are likely to be number of citations, share of global markets, winning international AI R&D competitions and contracts, and a prominent role in structuring AI governance, regulations, and standards. However, without specific stated goals set by the Zhiyuan Plan, it will be difficult to tell if its approach is or will be a success.

Beijing Academy of Artificial Intelligence

From the outset of the Zhiyuan Plan, BAAI has played a central role in establishing the complementary institutions and initiatives mentioned above. For its part, BAAI is to serve as an experimental hub for cooperation between the academic and corporate sectors while also receiving support from the government including funding and government data.

BAAI has announced four research focuses and houses two centers. The four research focuses are: mathematical foundations for AI, machine learning, intelligent system structures and chips, and intelligent information search and extraction. Its Center for Open Data Research is led by Tang Jie, an AI expert and Tsinghua University computer science professor. The Center's priorities are listed as: establishing open data standards, researching core open data technologies, and building an open data shared platform. Zeng Yi, one of China's most internationally engaged AI experts, leads the Research Center for Artificial

Intelligence Ethics and Safety. That center's work is to involve the exploration of ethics and safety or security principles, as well as the concrete implementation of those principles.

In May 2019, BAAI's Center for AI Ethics and Safety released the Beijing AI Principles covering research and development, governance, and usage of AI. The Principles are among the first guidelines for AI use in regards to privacy, freedom, and rights, raised by a Chinese government–backed organization. They have the potential to guide future regulation and shape Chinese discourse on new processes for creating rules and standards for emerging technologies.This new approach to AI development aims to move in lockstep with the advancements of leading AI companies and platforms, which emerge as regulatory stakeholders in the process.

BAAI's strategy appears to be: Provide space, provide funding, and let the experts handle the rest, while setting certain parameters for areas of research.

Finally, one of BAAI's most prominent campaigns to date has been the initiation and implementation of the Beijing Zhiyuan Scholars Plan, which is to induct 300 of the top Chinese AI experts over a period of three years. Scholars receive a generous stipend, as well as discretionary financial and data support to carry out their research. BAAI dean Huang Tiejun also explained that beyond this support, the recipients have the freedom to pursue whatever AI project they choose, except those appearing on a non-public negative list.

To this point, BAAI has announced the acceptance of 65 experts to the program: 47 established and 18 rising AI experts. Only four currently work in the private sector, while all others are professors, researchers, and leaders of academic and research institutions. Their research and program focuses include a vast range of topics, including fundamental principles of machine learning, small-sample machine learning methods, and deep neural networks.

BAAI's strategy appears to be: Provide space, provide funding, and let the experts handle the rest, while setting certain parameters for areas of research. If success in AI fields is to come from experts unleashed from the pressures of the market, this would seem an ideal environment for building a strong domestic ecosystem and leadership abroad. It is yet to be seen whether this government funding model, and the selection of individual researchers, will bear fruit.

Beijing's Efforts as an Example

The Zhiyuan Plan and BAAI are emblematic of several important trends that also appear in AI development initiatives across China.

The Zhiyuan Plan represents one model for responding to broad central government guidance plans. Typically, national development plans released by the State Council point out a general direction for economic and political priorities but

often lack detail regarding concrete steps. Central government departments and local governments are therefore left to determine the concrete measures to advance broader goals. The implementation of the Zhiyuan Plan is an example of the Chinese bureaucracy's recurring pattern of top-down directives and bottom-up execution.

The Zhiyuan Plan serves as a model for AI ecosystem development that other municipalities could emulate. Beijing's AI ecosystem is unique. It has one of the friendliest innovation regulatory environments and some of the world's most advanced companies and institutions in AI R&D. Not all cities boast the same resources, but other municipalities are showing signs of incorporating elements of the Zhiyuan Plan to advance the national AI ecosystem. Already, China's second AI Pilot Zone opened in Shanghai in May 2019. In September, plans were announced for another 18 national pilot zones by 2023. Not all municipal strategies and actions will be the same. They will vary based on each city's perceived advantages and individual officials' choices. For example, the annual World AI Conference held in Shanghai is quickly becoming a prominent forum for global AI innovation exhibition and debate. Organized by the Ministry of Science and Technology and the Shanghai municipal government, the conference was attended by over 170,000 people in its 2018 inaugural year, and in 2019 was headlined by Elon Musk and Jack Ma.

China's second AI Pilot Zone opened in Shanghai in May 2019. In September, plans were announced for another 18 national pilot zones by 2023.

The Zhiyuan Plan is intended to shore up weaknesses in China's AI ecosystem. Despite the rapid pace at which China's AI ecosystem is developing, there are still gaps between China's achievements and its ambitions. These include gathering and retaining top AI talent, building a strong discourse on ethics, governance, and technical standards, and developing a widely-used machine learning framework. BAAI is emblematic of the steps the Ministry of Science and Technology and Beijing municipal government are taking to shore up these areas given their critical role in furthering the State Council's AIDP. While elevating status and resources for engineers will help with talent, the other two gaps are far more complex. Open source framework development is a complex endeavor, and a framework is only truly successful if it is widely adopted. A strong discourse on ethics, governance, and standards would also require a careful balance of government incentives eliciting corporate buy-in, coupled with freedom for independent research and opinion.

The Beijing AI Principles are a potential map for future China AI regulation and governance. Among the goals AIDP set for China's domestic AI ecosystem and global footprint by 2025 one was the "initial establishment of AI laws and regulations, ethical norms and policy systems, and the formation of AI security assessment and control capabilities." Playing a central role in the development of international standards, governance practices, and regulations is one way to

help secure Chinese prominence in future markets. The Beijing AI Principles demonstrate China's focus on the topic and intention to collaborate internationally on AI governance and standards.

At the same time, the Beijing AI Principles likely have significant potential to guide future regulation of the field domestically. BAAI boasts collaboration among academics, the government, and industry. It therefore has the potential to work with companies to identify how to concretely work these still vague principles into actual AI research, development, and application. The future development of these or other principles could lay out a map for legal regulation, R&D, governance, and application of AI tools. ◆

Joseph Michaels contributed to this article.

Thomas Lehmann was previously the Development Manager of the Carnegie-Tsinghua Center for Global Policy (CTC), where he assisted in planning and implementation of the Technology and International Affairs Program. Tom is a member of the World Economic Forum's Beijing II Hub Global Shapers, serving as a co-founder of the Tech Forward Initiative. He is also an American Chamber of Commerce in China Leadership Development Fellow. Prior to joining CTC in 2016, Tom spent two years in a small mountainous village in rural Yunnan as a Teach for China Fellow working on projects addressing education inequality. Tom graduated from the University of California, Santa Barbara with a Bachelor's in Chinese Studies.

From All Sides

Four specialists describe their diverse approaches
To studying China's AI development

INTERVIEWS BY JOHANNA COSTIGAN

Like "artificial intelligence," a broad concept that engages numerous existing and so-far imagined technological, industrial, and social phenomena, the extended community of people around the world who study AI and the Chinese context is diverse. Specialists from a wide range of fields, previously focused on China or not, have found important events at play in the country's experience with machine learning, advanced automation, and data-driven technology.

In order to illuminate this diversity of interests at play, DigiChina reached out to four researchers who have engaged with China, policy, and AI from different perspectives. **Jeffrey Ding** has spent countless hours scouring, translating, and analyzing Chinese writings on AI and analyzing the intersection of nationality, technological reality, and politics. **Maya Wang** has published some of the world's most illuminating documentation of the ways AI can be employed in authoritarian politics and surveillance technology. **Paul Scharre** engages from a military and national security perspective, in which Chinese and U.S. military modernization efforts raise tough bilateral, international, and ethical dilemmas. Finally, **Danit Gal** is deeply engaged with Chinese and East Asian efforts to think through the ethics and governance of AI, observing how national and institutional factors play out in public and industry discourse.

The writer and scholar Johanna Costigan interviewed all four, and their conversations were edited for length and readability. –Ed.

JEFFREY DING
University of Oxford

Jeffery Ding is a Rhodes Scholar and Ph.D. candidate in International Relations at Oxford, but might be better known for his "(sometimes) weekly" ChinAI Newsletter featuring translations of Chinese-language tech- and AI-centered texts.

Describe your research focus related to China and AI.

I have always been interested in U.S.–China relations. I was born in Shanghai, and moved to Iowa City when I was three, and got interested in the relationship through high school debate and IR issues. At Oxford, I kept coming across documents in Chinese media that were just not being picked up on. I realized that there's a huge gap in translating documents; DigiChina is trying to fill that gap, but it also extends into the informal analysis. And who are the DigiChina equivalents in China?

To what extent do you think market forces in China will continue to push forward AI in a way they couldn't with biotech, and what consequences does that have for the CCP?

There are some areas where it might be justified to at least consider government intervention. Europe has no social media giants, because U.S. tech firms just dominated, and it's very hard to displace the cumulative gains of the leader. In that sense you could say the Great Firewall's economic protectionism was actually key to allowing China to have a competitive social media industry. You don't want the market to completely dominate, because especially in the global market, the existing leaders will continue to prevail.

Do you think other researchers or commentators sometimes miss the data privacy debate within China? For example, we've heard a lot of people saying Chinese people don't care about privacy.

Since I've been covering it for the past two years, there's definitely been a trend of more discussions about data privacy. Polls have come out saying that the great majority of Chinese internet users are concerned about AI and privacy. Some of it is lost in translation in the sense that the concept of privacy is pretty malleable and can mean different things. In the Chinese context, privacy protections are solely viewed in terms of data security, meaning companies don't lose your data. I do think there's a tendency in the West, because China is a place where censorship reigns and you have an authoritarian government, to just think that Chinese people are willing conspirators. I do think there's a dehumanizing component to this rhetoric.

How does China's AI approach compare to other leaders in the field?

My core argument is that no one does AI evaluation well, because national AI capability is such a fuzzy concept. A slice of Beijing where it's super high-tech might be much more advanced than a slice of Iowa. I looked at input and output —patents, publications, talent numbers. We also have to look at different aspects of the AI value chain. Sometimes we only see the sexy product applications, but there's also the technology layer and the foundational layer; it's Google and Facebook building fundamental architectures. When comparing different countries' AI abilities, it's probably more useful to clearly specify what you're trying to compare.

Are there any particular Chinese texts that people should pay attention to?

Probably 50% of what I translate are these new science and tech media platforms that mostly push their articles out on WeChat. There are about 10 of these that cover AI as one of their main areas of focus, so that's definitely a trend that I think more people should be aware of. ◆

MAYA WANG
Human Rights Watch

Maya Wang is a China Senior Researcher at Human Rights Watch, where she investigates issues including China's social credit system, protests, surveillance, and more. She is currently based in Hong Kong.

How would you describe your research focus?

I cover a range of human rights issues, from the use of torture to Xinjiang to Hong Kong to mass surveillance. We focus on different areas responding to the situation on the ground. Three or four years ago, I started to be informed about the social credit system by activists. That threat remains in the background, and I'm still interested, but other means of surveillance were also present in Xinjiang forming an interest in mass surveillance.

Would you say the extreme methods described in your HRW report "China's Algorithms of Repression" are indicative of the CCP's paranoia toward losing control of China's people, a particular bias directed at Uyghurs based on a combination of discrimination and reactions to the riots in 2008, or evidence that the government hopes to expand these practices beyond Xinjiang and is using Uyghurs there as a particularly extreme test case?

A bit of a combination. The use of low-tech mass surveillance has been a part of the CCP since it was a party. They set up systems like *hukou* [tying people's privileges and obligations to a hard-to-change locality of registration], *danwei* [work units that can shape far more than a person's job], and the *dang'an* political file system. These were old fashioned ways of controlling people, and when the party transitioned to a market-based economy in 1979, it quickly realized it posed a problem, because people were working in private companies which they have no control of.

So the rhetoric becomes extreme after that point, and they started using technology to augment control. The mechanisms were built over time, but the motivation was very similar. In 2000 the Chinese government enlisted foreigners' help to instate the Golden Shield project [also known as the Great Firewall].

Western critics often describe the crisis in Xinjiang as an instance in which the Chinese government is using surveillance tactics "against its own people." What nuance would you add to that assessment given the ethnic and cultural distinctiveness of Uyghur people? And what are some troubling examples of surveillance technologies targeted at all Chinese citizens regardless of ethnicity?

First of all, Xinjiang is an important example of how the human future could possibly look. It's not limited to that part of China or even China itself. You already see the collection of biometrics being used in other countries, including and in particular in the United States, where laws have not caught up with the technology. That collection is being centralized and used in violation of human rights, particularly the most vulnerable populations. Recent news would suggest the Trump administration wants to target immigrants via collection of DNA, and through big data, to very invasively trace people's movements.

In Xinjiang, the targeting of minorities and then the spread of these methodologies to the majority is concerning for all of us. The way the CCP is targeting Xinjiang is offering a new model of social control. It is not a one-size-fits all model, which the CCP and all oppressive governments tend to do. The collection of biometrics and real time monitoring, while allowing some life to happen, ensures that there is a greater system of punishment and reward, to make sure that those who are thinking undesirable thoughts against the government are controlled in a more extreme manner.

In the rest of China, we have documented the "police cloud," which has some similarities to the [Xinjiang-focused] Integrated Joint Operations Program program, though they are less intrusive. The cloud also tracks and predicts dissent and

involves the mass collection of DNA of ordinary people not connected to crimes and other biometrics.

Is there any validity in Chinese officials' demands that the United States should stop "interfering with the internal affairs of China"? Even if we could get it back, would a highly involved America be the best approach to curb these behaviors?

International human rights standards are for all governments and all people everywhere. The Chinese Constitution protects religious freedoms and expression. The Chinese government's argument that criticism of human rights is interference in domestic policy is one of those convenient arguments used to silence the world's

criticisms. They have no validity at all, and what's more concerning is no governments are taking severe actions against what's happening there. The U.S. recently blacklisted 28 entities that are contributing to human rights abuses, but they need to do more through implementations of the Magnitsky Act. We have too little leadership in the world today standing up for these rights.

Are there any particular Chinese texts or sources that people should be paying attention to?

For my own research, I read a lot of police accounts and technology companies' accounts on Weibo or WeChat. I read a lot of government documents as well, and there is a lot of material they just put online; a lot of information is publicly available. ◆

PAUL SCHARRE
Center for a New American Security

Paul Scharre is a Senior Fellow and Director of the Technology and National Security Program at the Center for a New American Security. Previously, he worked in the Office of the Secretary of Defense, where he played a key role in establishing policies on autonomous systems and emerging weapons technologies. He is the author of Army of None: Autonomous Weapons and the Future of War.

Talk a little bit about your career path and how you got interested in AI.

I got interested in military robotics when I was overseas in Iraq. I remember a very clear moment when I came across an IED. We discovered it before it exploded, and we had an explosives team come out to take care of it. I was expecting someone in a big bomb suit like you'd see in the movie "The Hurt Locker," and instead it was a robot. I worked on military robotics issues at the Pentagon after I left the Army. And one question that kept coming up was the role of autonomy in weapons systems.

In your view, is the "centaur" approach (the ability to "successfully marry human and machine intelligence into joint cognitive systems" as you define it in a Fall 2018 *Foreign Policy* piece) ideal?

What is the best possible outcome of developing automated weapons?

It's an optimal way to combine the benefits of both human and machine decision-making, which have different attributes and advantages in different settings. One of the challenges is how do you build joint cognitive architectures that combine the benefits of the speed and reliability and precision of machines with the broad and more flexible capacities of humans? Humans can apply judgment and context, which machines can't do today.

There are advantages to using machine intelligence in warfare, from the standpoint of reducing humanitarian harm and complying with standards.

There's often an assumption that fully autonomous weapons with no human involvement would be better from a military effectiveness standpoint. But the most effective militaries will be those that combine human and machine decision-making on the battlefield. The challenge is that the optimal blend of the two is going to change over time. We do not appear at this time to be anywhere near the sci-fi ability of AI. There are good reasons to think we want humans involved in these decisions for quite some time.

In the same article, you write about the potential catastrophe of humans "ceding effective control over what happens in war" and compare it to the power of algorithms controlling the stock market. It seems clear that practitioners in the field of finance have been willing to take AI risks. Given the "arms race in speed" you point out, is your assessment that defense and military officials will be as willing?

I think stocks are an interesting comparison; it's competitive, there's an advantage in speed, and adversaries are not going to trade intel on how their algorithm functions. It's an important cautionary tale as militaries look at this technology. Automation introduces risk in novel ways because of its scaling effects, in the number of incidents that might occur as well as speed. You can have an accident that spirals out of control very badly that has a widespread effect in ways that are not possible with people. Human traders would not have been able to make all those mistakes as quickly.

I really do think that defense organizations underestimate the risks of accidents with their own systems and are not adequately prepared for thinking about emerging technologies that might have very dangerous consequences.

If machine learning requires environments that are more stable than war zones, how can we give machines the chance to learn? Should we?

As machine learning systems overall come out of research labs into society, there are all these incidents where they don't function well in the real world because the training data is not robust enough or doesn't accurately reflect the situations they are put in. Thankfully, war is very rare. This means we don't have extensive data sets on what war looks like. For militaries, it's like training a

sports team to play a game once in a generation where the rules are constantly changing and the consequences are life and death.

Giving up on machine learning altogether would give up on significant advantages both in reducing civilian harm and military effectiveness. Machines can make some decisions in the real world and we anticipate some failures, but the autonomy needs to be bounded so that those failures are not catastrophic. The military is an inherently hazardous environment.

In a recent *Foreign Affairs* article, you point out that China has already begun developing a system of digital authoritarianism, via facial recognition, predictive policing, and other methods. What is the connection between digital authoritarianism and autonomous weaponry? Could one beget or normalize the other?

It's conceivable that the technologies that would be matured through widespread surveillance like facial recognition could have dual-use applications in military settings, and that's of course troubling as well. But I'm far more concerned with how authoritarian regimes are directly using the technology, including the lack of a system of checks and balances in place to manage that use.

What other current events are relevant to China and AI development?

China released two position papers at the UN meeting on lethal autonomous weapons and they basically said they endorse a treaty prohibiting the use of lethal autonomous weapons, although not their research and development, which is a significant loophole.

The really interesting thing about it was that they laid out these five key attributes that describe what constitutes an autonomous weapon; there's nothing that would meet these definitions. There's been a lot of speculation about whether this is a genuine olive branch toward some kind of arms control or an exercise in lawfare—a strategy to use international law to constrain other actors. Of course, while China is doing this they're engaging in systematic human rights abuses using AI technology. So there is certainly a disconnect.

Are there any particular AI texts that people should pay attention to?

I think that Elsa Kania's report Battlefield Singularity regarding China and AI is the best thing to read on Chinese developments in military AI. ◆

DANIT GAL
Keio University

Danit Gal is from Israel, but these days she doesn't stay anywhere for too long. Danit reads, writes, and speaks about topics relating to AI and ethics, particularly in China and the rest of East Asia. She has degrees from the University of Oxford and Yenching Academy, and travels around the world engaging with people on tough questions surrounding responsible AI development and implementation.

Talk about how you got here, your path toward China and AI research, and your focus now.

I was headed towards a DPhil in cybersecurity at Oxford, and narrowly escaped it—packed my bags and moved to China for the Yenching Academy at Peking University. I learned Mandarin Chinese at Oxford during my master's at the OII [Oxford Internet Institute], but it was very clean Mandarin. It took me a good while after getting to Beijing to understand the *er* [a key sound in the capital's local accent -Ed.] worked like magic.

I was involved with Tencent from the start, and in Beijing I started mingling with other companies like Baidu, Alibaba, and others. Being affiliated with both Peking and Tsinghua universities was a valuable asset that allowed me to reach out to and engage many companies. AI was the natural trajectory since everyone was so excited about it, and 2016 was a good time to dive straight into it before the market exploded with hype. Right now, I'm focused on understanding how AI ethics and governance play out on the national level of key AI actors, and to that end I'm trying to better familiarize myself with the complex landscape of countries I consider key for the future of AI.

Your work centers on the relationship between a country's cultural context and AI. In the case of China, what is one cultural misunderstanding you've observed that leads to an inaccurate assessment of AI policy or planning?

When people read the New Generation Artificial Intelligence Development Plan (AIDP), they have a reaction like "this is so terrifying, China wants to control the world." This is often linked-up with Putin's mis-contextualized comment on how those who control AI will control the world. Maybe—maybe not. In some ways I can see how this policy is perceived as threatening, but the Chinese researchers and practitioners I've engaged with have a very different perspective. Some of them seek to develop consciousness simply because they think AI without consciousness is more dangerous, because it doesn't care about humans, not because they believe this will make them superior to others. People tend to take China's government policy at face value with a devilish spin because it serves their interests.

In your recent paper "Perspectives and Approaches in AI Ethics: East Asia" you put China's attitude toward AI, particularly robots, in context with its East Asian neighbors. (AI as a partner in Japan, as a tool in South Korea, and China in-between.) What distinguishes these attitudes and the role that traditions play in forming that conception?

An example of the "tool" categorization would be Google Maps. It's purely functional. The partner view leans more toward AI that has a humanoid voice or appearance, assumes any kind of human attribute that makes itself more accessible and approachable. This ends up blurring that line. You could call Siri functional AI if it wasn't the

most sexually harassed "woman" in the world. That's why this spectrum is important, we need to understand how people use AI that was designed to be used as a tool but with the interface of a partner.

South Korea has an interesting approach which I think is spot-on: divide responsibilities between users, developers, and providers. Everyone gets a share of rights and obligations—it's a human team effort. Japan has a long-standing heritage of partnership with technology, which informs society; The Japanese government's concept of Society 5.0 promotes co-development and co-existence with AI. China is in the middle of that spectrum.

Can you talk about how Buddhism has contributed to perceptions of AI?

Both China and Japan have a long-standing Buddhist heritage. Both countries have Buddhist robot-monks that help worshipers engage on a deeper level. The key contribution Buddhist beliefs had on the idea of AI as a partner is the belief that everything, living or not, has the potential to become the Buddha by being cultivated to reach enlightenment. It's one of the most ancient instances of techno-animistic beliefs. This lends inanimate objects a special place in society, and the popularity of robots or AI systems that are designed to resemble humans or animals greatly benefit from that.

On the tool scale you have South Korea with a very strong Confucian heritage. You have a very clear idea of where you belong and what you should aspire towards. This resonates very strongly with the idea of anti-social development —technology that does not interfere with society.

In Japan, Society 5.0 promotes co-development and co-existence. Its Shinto heritage also contributes to deeply-embedded techno-animistic views. When you talk to people about what inspired them, everyone will refer to the same religious influences and popular culture undercurrents. It's not that every single person will say, "I dream about having a partner in AI," but when you talk to them, they say, "Why not? I grew up with these ideas and they seem nice."

Do you believe AI development is an issue that requires global cooperation? Should there be global standards on research and implementation of AI?

There are global actors shaping up to do exactly that, but they have not been successful so far. We have so many actors on the scene claiming that they're starting international AI policy but not actually doing it. Even if they're inclusive by name, developed countries have the time and capacity, while developing countries still care about electricity and running water. Most of the time, these global actors have a hard time getting international representation, and even when they do, implementing a wish-list of principles locally proves much more challenging than reaching consensus with 40-plus people. Standards can be a tried-and-tested path for that, but they take time to develop on an international level. Part of my work is doing exactly that with IEEE P7009, and it's very challenging.

Are there any particular Chinese texts or sources that people should be paying attention to?

Both the East and the West have movies and shows depicting robots as potential love interests. I'd recommend those to observe how culture interacts with policy and practice. People tend to underestimate the importance of culture, but even policymakers are affected by it. If you grow up thinking AI is going to go well for humanity, you tend to be more optimistic about it – it's that simple. I'd also recommend paying attention to the gender dynamics. If you look at current development trends, you'll see that every guy gets a robot girl but not every girl gets a robot boy. Humanizing AI doesn't only make it accessible— and to some of us, scary—it also makes it inherently unequal. ◆

Johanna Costigan is a writer from New York. She is currently an MSc candidate in Contemporary Chinese Studies at the University of Oxford. After graduating Bard College in 2017, she spent two years working at NYU Shanghai as a Writing and Speaking Fellow. While at NYU, she completed a year-long poetry writing and translation project followed by a year of research into China's memorials of the Second Sino-Japanese War. She is a member of the Young China Watchers Global Editorial Team. Her work has appeared in ChinaFile, SupChina, Nüvoices, and elsewhere.

How AI Can Better Serve People With Disabilities in China

'Information accessibility' is already on the agenda,
But civil society offers more

BY SIODHBHRA PARKIN

Technological innovation has tremendous potential to both improve as well as complicate the lives of people with disabilities (PWDs). That is perhaps especially true in AI. From more conventional assistive technologies like mobility aids and prosthetic devices to cutting-edge examples such as visual recognition and real-time speech-to-text software, AI is slowly changing the landscape for PWDs and other marginalized groups as they seek to navigate environments that are only rarely designed with their needs in mind.

Even as AI improves the tech available for the conservatively calculated 85 million people in China with some form of disability, however, experts are concerned the government and private sector alike continue to fundamentally misunderstand or misrepresent disability issues in China more broadly, even when AI projects occasionally pay lip service to inclusivity. This is particularly troubling at the present moment, as government officials are just beginning to turn their attention to revising the 2007 Regulations on the Employment of Persons With Disabilities.

These concerns stem in part from the fact that civil society has until relatively recently been sidelined in the process of formulating and implementing formal policy initiatives impacting PWDs, including tech-focused policies and initiatives. On a very basic level, this includes fundamental issues such as defining "disability" and identifying precisely what sort of support—including, increasingly, AI-based technological support—is made available to China's PWD community.

Even ascertaining basic facts about the size of the population of PWDs in China is difficult; while official statistics offer the aforementioned 85 million figure,

experts and human rights groups question the overly narrow definition used by government-affiliated groups like the China Disabled Person's Federation (CDPF), and estimate the number to be as high as 200 million. Part of this disconnect is the result of the Chinese government's tendency to shut out the voices of experts from within this community.

Certainly this is the case at the international level. The United Nations has formally noted that China has effectively excluded domestic nongovernmental organizations (NGOs) from participating in assessments of China's performance under the UN Convention on the Rights of Persons with Disabilities (CRPD), which the government ratified in 2008. In its initial review of China's progress in implementing the CRPD, the UN further noted that Chinese laws and policies are premised on the 'medical model' of conceptualizing disability rather than the more progressive human rights–based foundation upon which the CRPD was built—that is, official policy has a tendency to address disability as a problem to be mitigated rather than defining, valuing, and actively protecting the intrinsic rights of people with disabilities.

A 'Tag-On' Part for People With Disabilities in China's AI Economy

This fundamental miscasting of the role of the state, and subsequently, tech companies, address the needs of PWDs. For example, around the same time the State Council Office of Disabled Workers released new data indicating the unemployment rate of PWDs aged 16–59 was around 45% in 2018, Chinese tech giant JD.com (Jingdong) launched a new employment initiative that has been specifically marketed to PWDs called "Jingdong Micro Labor." The initiative offers relatively undemanding but stable jobs that provide a small amount of money in exchange for "tagging" or adding classification markers to various data within Jingdong's database.

The work requires little skill and offers few opportunities for advancement, though it is widely advertised as an AI-sector position. While that in itself is not inherently a problem—and in a society where there is widespread ableist discrimination, many PWDs welcome any employment opportunities at all— advocates are concerned the solution is rooted in a fundamental misconception of the rights of PWDs to attain fulfilling employment and social acceptance.

The State Council's 2017 New Generation Artificial Intelligence Development Plan makes no mention of using technology to expand available support for PWDs, but it does include a chapter on "Intelligent Health and Elder Care Systems" further suggesting that this approach is entrenched at the policy level. Disability is perceived as a range of medical problems to be treated rather than a starting point in a much larger conversation about how minority groups in China are enabled to engage society on their own terms.

Civil Society Offers a Fuller Vision

As a result, disability advocacy groups in China have worked hard to promote alternative versions of how AI, and technology more generally, can be shaped to actively affirm the rights of PWDs rather than simply mitigating the practical impact of their disabilities.

Groups like the One Plus One Disabled Persons Public Interest Group, a prominent disabled persons' organization (DPO), have led the charge in conducting public education campaigns and legal advocacy efforts aimed at pushing government agencies to realize the full extent and spirit of their commitments to the rights of PWDs more generally. One Plus One, along with other DPOs, has also sought to convince tech companies to invest more heavily in assistive technologies, not solely because they have the potential to transform the lives of users with disabilities, but in order to produce more diverse tools and software that have the potential to evolve into everyday, universal use. Speech-to-text software, for instance, was originally developed as an assistive technology for people with visual impairments and is now a key component of voice command functions on many types of smart devices.

Apart from DPOs, other groups with closer ties to industry and government have also emerged, pushing various types of reform. Seizing on the concept of "information accessibility," for example, the China Academy of Information and Communications Technology (CAICT), which is closely associated with the Ministry of Industry and Information Technology, and the Accessibility Research Association headquartered in Shenzhen have started publishing annual white papers assessing the current state of information accessibility in China and outlining steps to move the issue forward. The 2019 white paper encourages relevant government agencies to establish clearer standards for the private sector to meet as it researches and develops new technologies, including AI applications, that ensure "any person in any situation can equally and conveniently access information and use information," regardless of their age or disability status.

> *DPOs have sought to convince tech companies to invest more heavily in assistive technologies, not solely because they could transform the lives of users with disabilities, but in order to produce more diverse tools and software that could evolve into everyday, universal use.*

These efforts have not gone unnoticed by high-level government actors. In September 2018, Xi Jinping, Li Keqiang, and other key members of top Chinese leadership attended a conference on innovation and disability organized by a quasi-government agency, the Chinese Disabled Persons' Federation.

In his main address, Han Zheng, a vice premier of the State Council and member of the Party Politburo Standing Committee, urged his colleagues to "quicken the pace of establishing an environment of information accessibility,

conscientiously raise the level of technology use among people with disabilities, advance the safeguarding of the rights and interests of people with disabilities, promote the harmonious development of people with disabilities, and ensure people with disabilities share in the fruits of economic and social development."

If the government and the private sector make efforts to listen to the community Han's exhortation calls them to serve, there could be great potential to more holistically improve conditions for PWDs in China at a time when technological development is reordering so many aspects of everyday life. ◆

Siodhbhra Parkin is a China Program Manager at the Global Network for Public Interest Law, where she specializes in civil society engagement and legal advocacy issues. She was previously a Fellow at the Yale Law School Paul Tsai China Center, where she worked on domestic violence and LGBT rights programming. Before that, she was a Program Officer at the American Bar Association Rule of Law Initiative in Beijing. Parkin holds advanced degrees from Harvard University, the London School of Economics and Political Science, and the Renmin University of China. She earned her bachelor's degree at Harvard University. She is also the Director of the Serica Initiative, the newly established nonprofit arm of SupChina.

Planning Towards Policy Paralysis

*How government plans, combined with political tightening,
Form a barrier to AI governance*

BY KARMAN LUCERO

As President Xi Jinping and the Chinese Communist Party (CCP) exert more centralized and ideological control over legal institutions, the challenges of AI deployment across multiple industries and throughout society demand flexible and innovative responses. However, while the government's top-level plans for AI advancement call for policy adaptations and taking a lead in global regulation, the same plans appear to be shrinking the space for policy innovation even further.

AI, with its broad applications and vague definitions, is proving challenging for legal regimes across the world. In China, the complex dynamics of regulating AI coincide with the CCP's increased institutional and ideological control over legal institutions and the private sector. This combination is already having negative impacts on the Chinese legal system, particularly its capacity to respond to and regulate AI, because it affects the capacity of China's institutions to develop and govern.

National Ambitions for AI Amidst Increasingly Centralized Governance

When it comes to understanding China's bold AI-related declarations and actions, it is important to put them into this institutional context—to look beyond China's stated ambitions into the more nuanced reality of how "AI" is being described and used within China's political and legal institutions.

The Chinese system is defined in part by its political and legal centralization. Provincial and local governments, for example, do not pass legislation but rather "implement" laws passed by the National People's Congress and regulations issued by the State Council. The CCP also utilizes a political ideology that emphasizes its own singular legitimacy and wisdom to govern China. This ideology does not exist only at the top, but rather spreads throughout the various

bureaucratic and legal institutions across China. One routine characteristic of this system is the use of overarching plans to drive industrial and other important policy goals, and the headline-catching 2017 New Generation Artificial Intelligence Development Plan (AIDP) and its various local iterations continue this longstanding governance model.

The Chinese system in recent years grew even more centralized, as the CCP has worked to entrench its formal powers over state institutions, including government agencies and courts, as well as greater Chinese society. For example, a growing number of Party, regulatory, and court documents emphasize the "absolute leadership of the Party." The CCP has also established a number of new Party organizations that are outside the formal state hierarchy and are therefore effectively "extra-legal," or beyond the control or supervision of law, including the National Security Commission, which reports directly to Xi. The Cyberspace Administration of China, for its part, reports directly to the Party Central Committee.

Government agencies have reportedly responded to increased centralization and ideological control with fear and paralysis. When the correct way forward is unclear, sometimes it seems safer to do nothing at all.

Private companies, particularly tech companies, are also facing increased CCP interference and control. Tactics from buying company shares to requiring the establishment of Party Committees have, in the view of some analysts, allowed the CCP to "quasi-nationalize" private tech companies, transforming them into "state-overseen" enterprises.

This environment of increasing extralegal powers and personalized authority has exacerbated a bureaucratic paradigm that prioritizes political performance and loyalty, even over efficiency. Government agencies have reportedly responded to increased centralization and ideological control with fear and paralysis; when the correct way forward is unclear, sometimes it seems safer to do nothing at all.

China's plans and stated ambitions for the future of AI are far from exempt from these trends of centralization and political discipline.

Chinese Governance of AI and Its Effect on Politics and Development

At first glance, China's approach to the governance of AI appears similar to other countries. Other nations, as well as the OECD, have released similar "AI plans" and documents discussing the importance of ethics and principles when it comes to developing and deploying AI.

The difference is in the institutional details. Currently, the real substance of how AI is being governed across the world is not as much in the grand plans and pronouncements but rather in the particulars of how institutions and individuals affect the role of AI in their lives and communities. It is in this context that AI

appears to be revealing and potentially exacerbating shortcomings within China's political and legal institutions.

For example, while China's 2017 AI plan is no more vague than any other national document discussing AI, it signals not only intent, but political control. It is as much an announcement to the world that China will lead in AI as it is to domestic institutions that the Party will rule AI and the future it is to power.

If you combine the Party's assertion of control over AI with its tightened ideological control overall, as well as the indeterminate breadth of AI as a concept, such signaling could well exacerbate problems within government institutions. Overemphasis on "controlling" AI and/or "winning" the "AI race" could put further pressure China's institutions and reduce their regulatory flexibility.

Old-fashioned bureaucratic in-fighting could also stifle government innovation. There were at least fifteen central government agencies involved in drafting the AIDP. At the same time, Chinese government agencies have a history of infighting and competition. Given the complexity of AI as a legal concept and the political impetus to "win" at governing it, how are those institutions supposed to cooperate? Assigning responsibility to committees does not automatically lead to institution building.

Recourse to existing rules won't cut it, either. There are some laws on the books that govern the use of algorithms. However, many of these laws include idealistic, politically correct language that is difficult to implement. The 2017 Cybersecurity Law, for example, requires that network operators "respect social morality… and bear social responsibility." As the People's Daily describes, such language is having trouble shaping behavior in practice.

There is also a general dearth of regulation in a number of industries in which AI is being deployed. In transportation, for example, there is no national law that regulates safety or other key issues related to autonomous vehicles (though there are notices requiring licenses for smart maps in such vehicles and for "internet enabled" cars, which appear to be cars with some internet-accessing features). Local regulations too are lacking (especially compared to the United States, where 40 states have enacted legislation and/or executive orders). "Smart" medical products or mobile medical apps are also largely unregulated so far.

Paths Not Taken and Not Available for AI Governance in China

It is of course impossible to definitely say why there is a dearth of laws. It is possible, however, that the complexity of AI requires regulatory flexibility and institution building and that, currently, the CCP is placing heavy emphasis on increased bureaucratic and ideological control and the expense of flexibility. The desire for more control does not automatically translate into institution

building. In the current climate of institutional paralysis, government actors might not have incentives to make potentially risky legal innovations, and they might instead continue to stagnate.

While local governments in China have a history of innovation in certain contexts, this dynamic appears to be declining. One unsigned commentary laments that, while information technology was supposed to make life easier for bureaucrats, it appears to instead have added more hurdles and vectors for political risk. AI, being both complex in a way that requires innovation and politically important in a way that requires signaling, is straining individuals within China's bureaucracies.

One route taken in many countries around the world is less open for China. Since legislatures across the world have generally been slow to respond to the advent of AI, civil society organizations have played a large role in the nascent development of AI governance, since they can act spontaneously as a check on both government and private sector power. The AI Now Institute, for instance, has published several reports on different uses of algorithms and their impact on society. But President Xi has overseen a large crackdown on civil organizations, particularly law-oriented ones. As such, civil society organizations within China lack the ability to fill the governance and public interest gap left by a bureaucracy lagging the development of technology.

The CCP is currently closing legal spaces across the board while simultaneously emphasizing the importance of governing AI. There is some evidence that political signaling in the AI space is taking precedence over realistic institutional creativity.

Courts within China have recently been institutionally innovative, but innovations generally serve to insulate courts from criticism and political risk rather than increase their authority or capacity to address complex issues. With AI's breadth and political import, Chinese courts might not only face unprecedented legal challenges in cases that involve AI, they might also face unprecedented political pressure to avoid any chance of hindering the CCP's plans for AI. The courts are still worth watching, however: The Supreme People's Court might issue guidance for lower courts in terms of dealing with cases that involve algorithms. If such guidance is issued, it will be important to pay attention to local courts and how they react to cases involving algorithmic decision making.

AI Governance Dilemmas Have Broader Political Effects

The complexity and novelty of governing AI requires space for regulatory flexibility and compliance—the space to make mistakes and experiment. The CCP is currently closing legal spaces across the board while simultaneously emphasizing the importance of governing AI successfully, and there is some evidence that political signaling in the AI space is taking precedence over realistic institutional creativity.

More national drives for local governments to fund AI winners mean more money out the door, and many local governments in China already face potentially unsustainable levels of debt. AI projects are far from guaranteed to break this cycle, as many of the government-funded "AI startups" do not make much use of AI, and many suffer from unsustainable business models.

To understand China's future AI governance and technological development, one must go beyond the stated principles and ambitions and observe the development (or lack thereof) of institutions. There are signs that the CCP recognizes problems of bureaucratic inaction and is looking for ways to improve local institutional capacity. Whether efforts to meet these challenges will prove successful remains an open question. ◆

Karman Lucero is currently a Fellow at the Paul Tsai China Center at Yale Law School, where he works on issues related to the intersections of technology, law, and U.S.–China relations and other issues. Before joining the Center, he worked at the Data & Society Research Institute on its Intelligence and Autonomy project, conducting research on the impact of artificial intelligence on law and society. Previously, he worked for Microsoft's U.S. Government Affairs Office, where he focused on issues related to artificial intelligence law and policy, telecom law and policy, and criminal justice reform. He received a J.D. from Columbia Law School, where he was a Harlan Fiske Stone Scholar and Editor-In-Chief of the Columbia Journal of Asian Law. Prior to law school, he was a Teaching Fellow in Manghuai Township in rural Yunnan with Teach for China. Karman speaks and reads Mandarin Chinese.