

MCKINSEY GLOBAL INSTITUTE

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR CHINA

APRIL 2017

DISCUSSION PAPER
PRESENTED AT THE 2017 CHINA DEVELOPMENT FORUM



Dominic Barton | London
Jonathan Woetzel | Shanghai
Jeongmin Seong | Shanghai
Qinzheng Tian | Beijing



MCKINSEY
GLOBAL
INSTITUTE
RESEARCH.
INSIGHT.
IMPACT.

MCKINSEY GLOBAL INSTITUTE

Since its founding in 1990, the McKinsey Global Institute (MGI) has sought to develop a deeper understanding of the evolving global economy. As the business and economics research arm of McKinsey & Company, MGI aims to provide leaders in the commercial, public, and social sectors with the facts and insights on which to base management and policy decisions. For the second year running, the Lauder Institute at the University of Pennsylvania ranked MGI the world's number-one private-sector think tank in its 2016 Global Think Tank Index.

MGI research combines the disciplines of economics and management, employing the analytical tools of economics with the insights of business leaders. Our “micro-to-macro” methodology examines microeconomic industry trends to better understand the broad macroeconomic forces affecting business strategy and public policy. MGI's in-depth reports have covered more than 20 countries and 30 industries. Current research focuses on six themes: productivity and growth, natural resources, labor markets, the evolution of global financial markets, the economic impact of technology and innovation, and urbanization.

Recent reports have assessed the economic benefits of tackling gender inequality, a new era of global competition, Chinese innovation, and digital globalization. MGI is led by four McKinsey & Company senior partners: Jacques Bughin, James Manyika, Jonathan Woetzel, and Frank Mattern, MGI's chairman. Michael Chui, Susan Lund, Anu Madgavkar, Sree Ramaswamy, and Jaana Remes serve as MGI partners. Project teams are led by the MGI partners and a group of senior fellows, and include consultants from McKinsey offices around the world. These teams draw on McKinsey's global network of partners and industry and management experts. Input is provided by the MGI Council, which coleads projects and provides guidance; members are Andres Cadena, Sandrine Devillard, Richard Dobbs, Katy George, Rajat Gupta, Eric Hazan, Eric Labaye, Acha Leke, Scott Nyquist, Gary Pinkus, Oliver Tonby, and Eckart Windhagen. In addition, leading economists, including Nobel laureates, act as research advisers. The partners of McKinsey fund MGI's research; it is not commissioned by any business, government, or other institution. For further information about MGI and to download reports, please visit www.mckinsey.com/mgi.

IN BRIEF

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR CHINA

Artificial intelligence, or the idea that computer systems can perform functions typically associated with the human mind, has gone from futuristic speculation to present-day reality. When the AlphaGo computer program defeated Lee Sedol, a nine-dan professional master, at the game of Go in 2016, it signaled to the world that it is indeed possible for machines to think a bit like humans—and even exceed their capabilities.

Thanks to advances in data collection and aggregation, algorithms, and processing power, computer scientists have achieved significant breakthroughs in artificial intelligence. Where computer systems once had to be programmed to execute rigidly defined tasks, they can now be given a generalized strategy for learning, enabling them to adapt to new data inputs without being explicitly reprogrammed. Today many machine learning systems have already been developed for commercial use. The applications are tremendously varied, and adoption is growing rapidly in sectors such as finance, health care, and manufacturing.

Because they can dramatically boost productivity, AI technologies may have a disruptive impact on China's economic growth and on its workforce. A McKinsey Global Institute study published earlier this year estimated that half of all work activities in China could be automated, making it the nation with the world's largest automation potential. Hundreds of millions of Chinese workers could be affected, and jobs made up of routine work activities and predictable, programmable tasks will be particularly vulnerable. While impact on the labor market is likely to be gradual at the aggregate level, it can be sudden and dramatic at the level of specific work activities, rendering some jobs obsolete fairly quickly. Overall, AI will raise the premium placed on digital skills while reducing demand for medium- and low-skill workers, potentially exacerbating income inequality.

On the plus side, AI's effect on productivity could be crucial to China's future economic growth as the population ages. According to a McKinsey Global Institute report, AI-led automation can give the Chinese economy a productivity injection that would add 0.8 to 1.4 percentage points to GDP growth annually, depending on the speed of adoption.

With its biggest tech companies driving significant investments in R&D, China is one of the leading global hubs of AI development. Its advantages include a vast population and diverse industry mix that have the potential to generate huge volumes of data and provide an enormous market. But China will need to focus on building its innovation capacity. The United States and the United Kingdom are currently producing more influential AI research, and the more robust US ecosystem nurtures more competitive AI startups. Realizing AI's economic potential in China also depends on its actual adoption—not just among the technology giants but across China's traditional industries. Achieving this goal will require building strategic awareness among business leaders, developing technical know-how, and overcoming implementation costs.

AI capabilities have exciting and far-reaching potential to enhance human welfare by improving health care, the environment, security, and education. At the same time, AI also raises complex ethical, legal, and security questions surrounding issues such as privacy, discrimination, liability, and regulation. Prudential governance should be put in place as AI is introduced into society on a broader scale.

Although the market will drive the development and adoption of AI, the right policy framework can establish a healthy environment for growth. Five priorities can form the basis of China's AI strategy: building a robust data ecosystem; spurring adoption of AI within traditional industries; strengthening the pipeline of specialized AI talent; ensuring that education and training systems are up to the challenge; and establishing an ethical and legal consensus among Chinese citizens and in the global community.

The technology industry is becoming increasingly global. China has the capability and opportunity to lead international collaboration in the development and governance of AI, ensuring that this breakthrough technology will positively contribute to the general welfare of all humanity.

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR CHINA

Artificial intelligence (AI), or the idea that computer systems can perform functions typically associated with the human mind, has suddenly gone from futuristic speculation to present-day reality. Computer scientists have achieved significant breakthroughs in machine learning and deep learning, giving machines cognitive and predictive capabilities. Today these systems are already being deployed in real-world situations.

A JOURNEY TO THE BRINK OF TRANSFORMATION

AI is defined as machines mimicking cognitive functions typically associated with the human mind. This concept has long been the stuff of speculation and science fiction—and the optimism surrounding it intensified after a number of initial theoretical advances were made in the 1950s and '60s. But that wave of momentum fizzled in the face of technical obstacles. With AI failing to meet expectations, the field experienced a long fallow period.¹ Subsequent decades brought a few successes (such as IBM's "Deep Blue" supercomputer defeating Gary Kasparov in chess), but the real-world use cases were too isolated to support mass commercialization.

Fast forward to the 21st century. Breakthroughs in data collection and aggregation, computing power, and algorithms (especially machine learning) enabled revolutionary technical advances. In one widely hailed milestone, Google's AlphaGo computer program defeated a human champion in the game of Go, which had been traditionally regarded as unsolvable by machines.

But the advances are not only happening at the theoretical frontiers of the field. Analytics tools utilizing machine learning are the precursors of tomorrow's super-intelligent systems, and many of them are already on the market. Adoption is growing rapidly in sectors such as finance, health care, and manufacturing. Global venture capital funding has grown from \$589 million in 2012 to over \$5 billion in 2016.² McKinsey estimates that the total market for AI applications will reach \$127 billion by 2025.

UNDERSTANDING AI AND WHAT IT CAN DO

Traditionally, we have used the processing power of computers to generate output more efficiently (for example, doing faster and more complex computations than humans can perform). Conventional software programs have always been coded with specific instructions on the tasks they need to execute.

AI systems take a very different approach. They can sift through enormous "big data" sets to find patterns, associations, and insights—and as they do, they employ a generalized strategy for learning. This enables them to adapt to new data inputs without being explicitly reprogrammed. Systems utilizing machine learning³ have induction and decision-making capabilities—and the systems being developed on the frontiers of this field push the

¹ Daniel Crevier, *AI: The tumultuous history of the search for artificial intelligence*, Basic Books, 1993.

² "The 2016 AI recap: Startups see record high in deals and funding," CB Insights blog, January 19, 2017, <https://www.cbinsights.com/blog/artificial-intelligence-startup-funding/>.

³ Machine learning is one of the most important technical developments in the field of AI. Building on the premise that the human cognitive process can be represented by mathematical models, it feeds a huge volume of data into an algorithm that is essentially a generalized strategy for learning. It then "trains" the machine to derive a rule or procedure for interpreting data or making predictions.

boundaries even further with deep learning.⁴ These computer systems can learn, discover, and apply rules by themselves.

While recent breakthroughs in deep learning have produced AI systems that can match or surpass human intelligence in certain key functions, we are still decades away from “general AI”—or machines that can perform the full range of cognitive tasks that humans can do. But many machine learning systems have already been developed for specific commercial uses, and the applications are tremendously varied. They can provide customer service, manage logistics, monitor equipment on factory floors, optimize energy consumption, and analyze medical records. Recent McKinsey Global Institute (MGI) research finds that machine learning techniques have wide applicability in virtually every industry.⁵

It is useful to think of AI capabilities in four main categories:

- **Perception** involves collecting and interpreting information to sense the world and describe it. These capabilities include natural language processing, computer vision, and audio processing.
- **Prediction** involves using reasoning to anticipate behaviors and results. Such technologies are used, for example, to develop precisely targeted advertising for specific customers.
- **Prescription** is principally concerned with what to do to achieve goals. It has a variety of use cases, including route planning, drug discovery, and dynamic pricing.
- Last but not least, AI can be combined with complementary technologies such as robotics to provide **integrated solutions**. These include autonomous driving, robotic surgery, and household robots that respond to stimuli.

The current degree of commercialization varies for each type of AI functionality. While systems with perceptive and predictive capabilities are already on the market, more prescriptive tools and integrated solutions are still under development (Exhibit 1).

AI'S FUTURE TRAJECTORY: PROFOUND CHALLENGES AND POSSIBILITIES

The technology advances of the past mainly enhanced capabilities to execute clearly delineated tasks in production. But now AI enables machines to react and adapt in order to optimize results. Together with technologies such as the Internet of Things (IoT) and robotics, it can create an integrated cyber-physical world.

Current momentum points to the likelihood that AI technologies will eventually be embraced globally in an even wider variety of settings and industries—and one of the most important consequences would be machines handling a variety of tasks that have always been performed by humans. An MGI report analyzed more than 2,000 work activities across 800 occupations in the global economy. Already it appears to be technically feasible that 50 percent of today's work activities could be automated using currently demonstrated technologies.

But technical feasibility is only one factor affecting the pace and extent of automation. Others include the cost of developing and deploying specific applications, labor market

⁴ Deep learning is a sub-branch of machine learning on the frontier of computer science. This technology involves software-based calculators that approximate the function of neurons in a brain; they are connected to form a hierarchical “neural network.” Instead of the “shallow” learning algorithms and manual feature extraction that characterize traditional machine learning, deep learning runs data inputs through multiple layers of non-linear processing units, automatically extracts data features, and uses the previous layer's output as an input for the next layer. The intricacy of these neural networks enables even more sophisticated capabilities, such as image recognition and natural language generation.

⁵ *The age of analytics: Competing in a data-driven world*, McKinsey Global Institute, December 2016.

dynamics, the economic benefits, and regulatory and social acceptance. Taking these factors into account, MGI's research on automation finds that it might take until 2055 for half of all current work activities to actually become automated—but there is a fair degree of uncertainty in this timing. In an aggressive adoption scenario, that level of automation could occur 20 years sooner, and in a late adoption scenario, it might occur 20 years later.⁶

Further down the road, AI could be a powerful tool to apply to some of society's central challenges. In health care, AI will greatly enhance our capability to analyze the human genome and develop personalized and more effective treatments for each patient. It could radically accelerate efforts to cure cancer, Alzheimer's, and other diseases. AI systems can analyze weather patterns and improve energy efficiency on a wide scale, enhancing our ability to monitor and combat climate change. And the possibilities are not even earthbound; AI systems could one day pioneer exploration of Mars and the outer reaches of space.

Exhibit 1

The current degree of commercialization varies across AI technologies

Functionality	Maturity	Use cases
Perception		IBM Software can review medical images with the same accuracy as human radiologists
		Iflytek Voice assistant app transcribes spoken Mandarin into text
Prediction		Netflix Algorithm suggests films and TV shows to customers based on their previous viewing history and ratings
		Capital One Algorithm predicts customers' purchasing behavior
Prescription		Wealthfront AI-driven platform provides automated advice to customers on asset allocation and wealth management
		Google AI can produce surrealistic "artwork" from white noise or images
Integrated solutions		Amazon Smart speaker devices can control home appliances
		Baidu Autonomous cars operate within known and limited environments

SOURCE: MIT Technology Review; TechNode; WealthManagement.com; Google Research blog; McKinsey Global Institute analysis

WHAT DOES AI MEAN FOR CHINA?

With its biggest tech companies driving momentum for R&D, China is one of the leading global hubs of AI development. Its vast population and diverse industry mix have the potential to generate huge volumes of data and provide an enormous market. Wide adoption of AI technologies could be crucial to China's future economic growth as the nation's population ages, heightening the need to accelerate productivity growth. Some of the required building blocks include a more open data environment and well-trained data science talent. But AI also poses complex social and economic questions that will require careful consideration.

CHINA'S POSITION IN AI DEVELOPMENT

China and the United States are currently the world leaders in AI development. In 2015 alone, they accounted for nearly 10,000 papers on AI published in academic journals, while the United Kingdom, India, Germany, and Japan combined to produce only about half as many scholarly research articles.⁷

⁶ This assumes that the human labor replaced by automation would rejoin the workforce and be as productive as it was in 2014. *A future that works: Automation, employment, and productivity*, McKinsey Global Institute, January 2017.

⁷ SCImago Journal & Country Rank, 2015

Much of the momentum behind AI in China is being driven by private-sector tech firms. Aided by huge volumes of search data and their many product lines, some of China’s Internet giants are on the cutting edge of technologies such as image and voice recognition.⁸ These capabilities have been integrated into new products, including automated personal assistants, autonomous cars, and so forth.

China has reason to feel optimistic about its role in a future defined by AI. Its huge population can generate a tremendous volume of data, which is a prerequisite for “training” AI systems. China also has the advantage of “economies of scope”: its wide range of industries provide a fertile market for deployment.

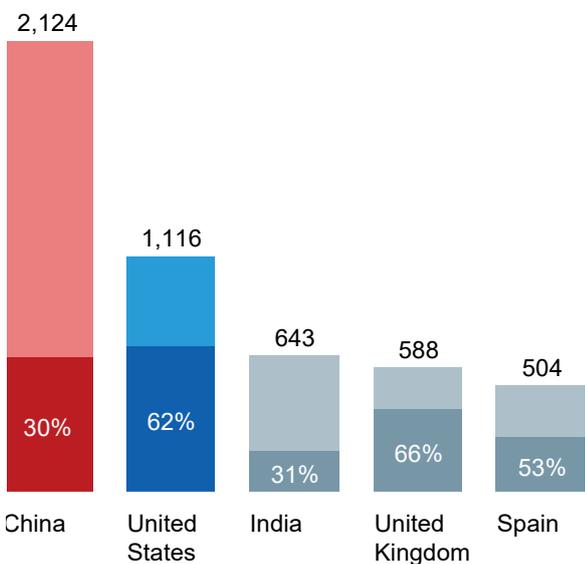
But it will take a sustained effort to stay at the forefront of such a rapidly evolving field and maximize the economic potential of these technologies. China will need to focus intently on bolstering its capacity for innovation. For example, while Chinese academics have actually published even more papers on AI than US researchers, their papers have not generated the same impact as those by US or UK authors (Exhibit 2).

Exhibit 2

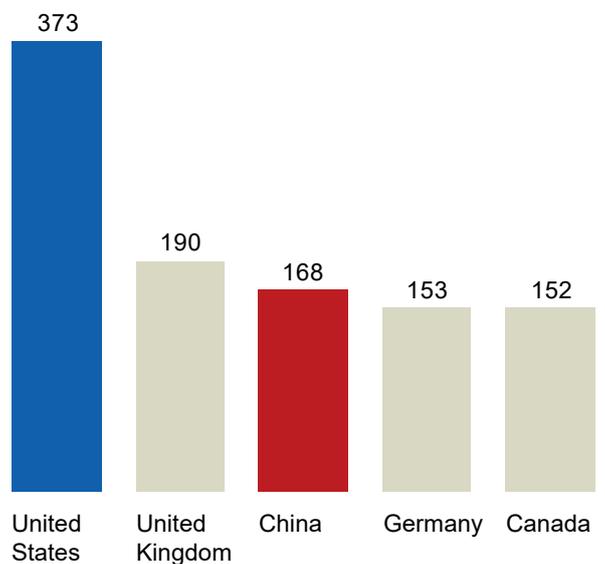
Although China produces a large number of widely cited AI-related papers, US and UK research remains more influential

While China ranks first for absolute AI citations, the United States holds an edge when self-citations are taken out

Number of AI publications cited ■ Self-citations¹
 ■ Other citations



Publication influence
 H-index²



1 Self-citation occurs when a journal cites another article published in the same journal.

2 The H-index ranks both the productivity of scholars and the citation impact of their publications. A higher H-index number indicates more publications that are widely cited.

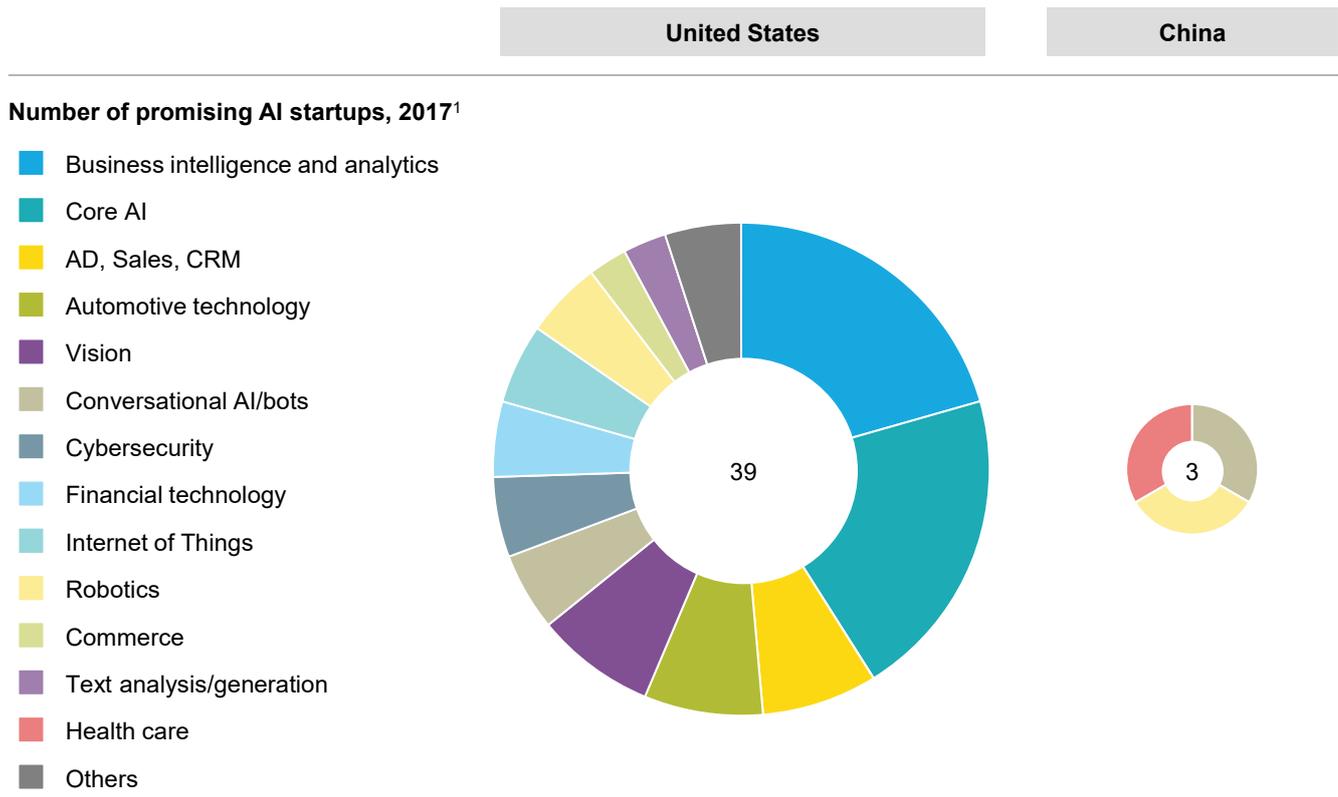
SOURCE: SCImago Journal Rank 2015; McKinsey Global Institute analysis

Furthermore, China does not yet have the same kind of vibrant AI ecosystem as the United States, which has produced substantially more AI startup companies than China (Exhibit 3). The US ecosystem is large, innovative, and diverse (including research institutions and universities as well as private companies). Building on all the well-established strengths of the Silicon Valley tech sector, it has considerable advantages that are difficult to replicate.

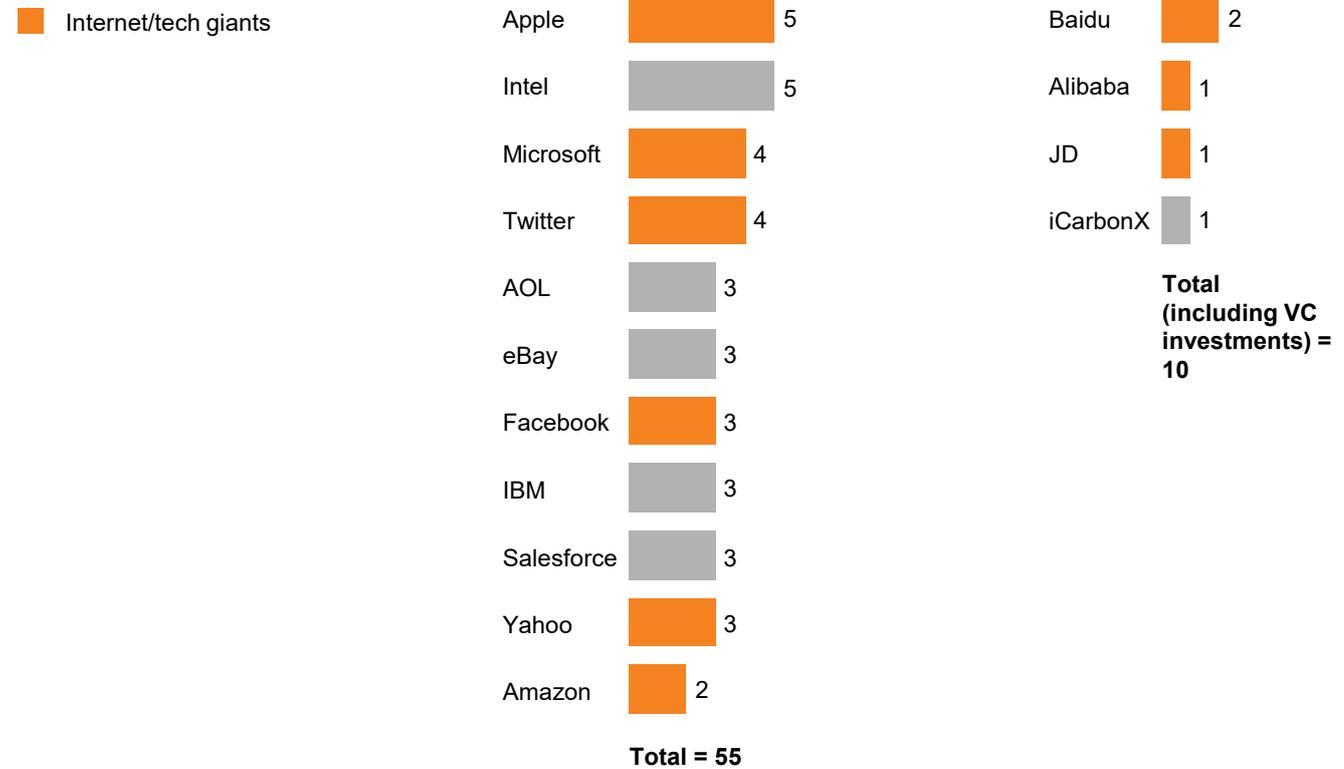
⁸ “Why deep learning is suddenly changing your life,” *Fortune*, September 28, 2016.

Exhibit 3

The United States has a more robust AI startup ecosystem than China



Major M&A deals, January 2011–February 2017



¹ Chinese and US companies identified from the 50 largest startups, ranked by total funds raised, from CB Insight's AI 100 list.

SOURCE: CB Insights; McKinsey analysis

It is helpful to view China's challenge through the lens of the three key building blocks of AI development: data, algorithms, and computing power.

Data

Just as humans are fueled by food, AI cannot run without a steady diet of data. These systems must be fed huge quantities of data so they can "train" themselves and continuously improve and refine their output. But several issues with data could stymie China's AI development.

First of all, while China's technology giants collect vast troves of data through their proprietary platforms, China is lagging behind the United States in creating a data-friendly ecosystem with unified standards and cross-platform sharing. Second, countries around the world have found that opening government data sets spurs private-sector innovation, but China has relatively little public-sector data accessible for exploration (Exhibit 4). And last but not least, limitations on cross-border data flows put China at a disadvantage for global collaboration.

Exhibit 4

China ranks 93rd globally for the openness of government data

	Global ranking for data openness, 2015 ¹	
	United States	China
Weather forecasts	13	80
Water quality	15	74
National statistics	1	106
Government spending	8	82
Government procurement	1	36
Location (postcodes)	49	61
National maps	1	88
Legislation	1	39
Land ownership	66	85
Pollutant emissions	1	30
Election results	83	87
Company register	33	71
Government budget	1	49
Overall	8	93

¹ The assessment in each data category considers 10 aspects that contribute to public accessibility. Among these criteria are whether the data is published online and whether it is free, up to date, and machine readable.

SOURCE: Open Knowledge International, 2015; McKinsey Global Institute analysis

Algorithms

At the application level, China is on a par with other countries in terms of algorithm development. In fact, Chinese players have achieved breakthroughs in developing AI algorithms used in voice recognition and targeted advertising. Thanks to global open-source platforms, Chinese companies are able to quickly replicate the most advanced algorithms developed anywhere in the world.

Nevertheless, China lags behind the United States and the United Kingdom in terms of fundamental research that advances the field of AI. One of the major reasons is simply a talent shortage. It will be critical for China to make a sustained push in AI talent development. Over half of the data scientists in the United States have more than 10 years of work experience, while up to 40 percent in China have less than five years of experience.⁹

China currently has fewer than 30 university research labs focused on AI, and they are not able to develop enough talent to meet the recruiting needs of China's AI industry. In addition, Chinese AI scientists have disproportionately specialized in areas such as computer vision and voice recognition, creating gaps in some other areas. Chinese university programs in AI could also benefit from imposing higher mathematical and statistical requirements and prioritizing efforts to stay current with global developments in the field. It is also worth considering whether the model for awarding research funding could be changed to promote greater innovation.

Computing power

Computing power is not an immediate bottleneck for the commercial development of AI in China. With microprocessors widely available on the global market, computing power has become a commodity that can be easily procured.

But China cannot afford to overlook the importance of developing its own advanced semiconductor, microprocessor, and high-performance computing technologies. Computing power is part of the basic infrastructure underlying AI and is thus of significant strategic importance.

China has historically been heavily dependent on foreign supplies for microchips. For some types of high-value semiconductors, China must rely on imports for virtually all of its needs.¹⁰ But in 2015, the US government banned Intel, Nvidia, and AMD, the three largest chip suppliers globally, from selling high-end supercomputer chips to the Chinese government.¹¹ Stronger control over the supply of core technologies can potentially improve China's future ability to deploy AI systems more widely.

To address this situation, the Chinese government published its National Guidelines for Development and Promotion of the Integrated Circuit (IC) Industry in 2014 and the "Made in China 2025" policy. The government also launched a national IC investment fund with more than \$20 billion raised so far. These initiatives are beginning to yield results: in June 2016, China unveiled Sunway TaihuLight, which broke records as the world's fastest supercomputer and contained no US-developed processors. The government's up-front investment is a long-term bet that should have a significant ripple effect as it encourages private companies to play an active role.

Specialized processors, such as graphics processing units that can carry out massive complex computations, are of particular importance for AI. As China grows its IC industry, it should also pay sufficient attention to the development of such types of processors.

As China considers its strategy to develop artificial intelligence, it is important to note that the technology industry is becoming increasingly global. All aspects of the AI value chain, from fundamental research to application development to hardware manufacturing, involve global collaboration. In addition to building its own data ecosystem, pipeline of data science

⁹ "Ceiling hit by search model, Baidu uplifts AI strategy with new CEO," 21st Century Business Herald, January 18, 2017, available at http://epaper.21jingji.com/html/2017-01/18/content_54928.htm

¹⁰ "Chips on their shoulders," *The Economist*, January 23, 2016.

¹¹ Khalid Moammer, "US government bans Intel, Nvidia, and AMD from selling high-end chips to the Chinese government," WCCFtech, April 14, 2015, <http://wccftech.com/us-government-bans-intel-nvidia-amd-chips-china/>.

research talent, and semiconductor industry, China needs to make sure that its AI industry is built on an open system that is integrated to the global market.

ECONOMIC IMPLICATIONS OF AI DEVELOPMENT

AI represents a significant opportunity for China to accelerate productivity growth, which is a crucial concern as the population ages. However, policy makers will also need to consider and prepare for the potential labor market disruptions it could unleash.

In the past few decades, China has benefited greatly from a “demographic dividend,” as its expanding labor force fueled economic growth. But China will lose that momentum as its population ages. The country’s working-age population has already peaked and will continue to shrink in the decades ahead. This demographic trend implies that China would fall well short of the workforce needed to sustain economic growth at current productivity levels. The only alternative for maintaining momentum would be to sharply accelerate productivity growth.¹²

AI could partially close this gap. AI systems can improve productivity by completing existing job activities more efficiently, either by assisting or replacing humans. Intel, for example, collects mass data alongside its chip manufacturing process, and the company used to rely on human employees to do root-cause analysis with the data if an error appeared. But now machine learning can complete this task much faster than humans; algorithms can sift through thousands of data points about each chip to find the common patterns among those with defects. Furthermore, AI can make industrial machinery, supply chains, logistics routing, and other processes more efficient. AI applications could create dramatic efficiencies by predicting failures, identifying bottlenecks, and automating processes and decisions.

A large share of the Chinese economy consists of accommodation and food services, manufacturing, agriculture, and other sectors in which a disproportionate share of the work consists of routine tasks that can be automated. According to an MGI report, AI-led automation can give the Chinese economy a productivity injection that would add 0.8 to 1.4 percentage points to GDP growth annually, depending on the speed of adoption.

In addition to improving productivity, the rise of AI is also likely to create new products and services, and hence new occupations and businesses. Just a few decades ago, no one would have imagined the many jobs now associated with the Internet economy, and AI could have a similarly transformative effect.

AI has the potential to dramatically boost productivity growth, but that may come at the expense of greater income inequality. Fewer people will be needed in roles such as customer service; consider, for instance, how Alibaba has incorporated AI customer support in its mobile payment app. Overall, AI will increase the trend of so-called “skill-biased technological change”—that is, there will be a new premium on digital skills but at the same time, there will be reduced demand for medium- and low-skill workers. This may decrease total labor demand, and while the average income may rise, the distribution may become even more skewed toward people with the right skill sets. The “digital divide” could manifest as a societal divide.

Overall, China has more labor associated with activities that can be technically automated than any other country in the world. MGI estimates that 51 percent of work activities in China can be automated, affecting the equivalent of 394 million full-time employees.¹³ However, even in an early adoption scenario, in which ~90% of work activities will be automated by 2055, China may still face a shortfall of the labor needed to meet its GDP

¹² *Global growth: Can productivity save the day in an aging world?* McKinsey Global Institute, January 2015.

¹³ *A future that works: Automation, employment, and productivity*, McKinsey Global Institute, January 2017.

growth target of 4-5 percent. This would leave the nation looking for additional levers to enhance productivity.

Jobs made up of routine work activities and predictable, programmable tasks will be particularly vulnerable to replacement by AI. Because of the cost-benefit calculation, middle-skill workers may bear the initial brunt, while lower-paying positions may survive longer. However, this is not to say that high-skill jobs today will be completely shielded from disruption. Many of the tasks performed by professionals with specialized knowledge and experience, such as doctors, may be subject to automation; these jobs could change to focus more heavily on personal interactions.¹⁴ Many jobs will not disappear, but their mix of activities will change, and education and training systems will need to respond.

A recent US government report looked at the types of AI-related jobs that could prevail in the future.¹⁵ It grouped them into four categories: engagement jobs requiring work alongside AI systems to complete complex tasks (such as nurses using AI applications to do routine patient checks); development jobs to create AI technologies and applications (such as data scientists and software developers); supervision jobs that monitor, license, or repair AI systems (such as technicians servicing AI robots); and jobs that respond to AI-driven paradigm shifts (such as lawyers creating legal frameworks around AI, or urban planners who create environments that can accommodate autonomous cars).

The combination of sharply increased demand for advanced digital skills and a growing surplus of lower-skill labor may exacerbate inequality. Some segments of the population may be particularly vulnerable. Women, for example, currently account for less than 20 percent of computer science graduates.¹⁶ They are over-represented in vocations made up of routine tasks that lend themselves to automation and under-represented in technology and management roles. China scored 83.8 for women's employment in the recent MasterCard Worldwide Index of Women's Advancement, but only 27.8 for women in leadership, indicating a high level of gender inequality in high-skill roles.¹⁷ AI may widen the gender gap.

Similarly, the growing adoption of AI may underscore existing divisions between the more prosperous coastal regions and less-developed inland regions as well as the opportunity gap between urban and rural areas. Taking these possibilities into account will be a vital part of planning for a future that includes a major role for AI systems.

SOCIAL IMPLICATIONS

AI capabilities have exciting and far-reaching potential to enhance human welfare by improving health care, the environment, security, and education. At the same time, as it blurs the lines between the physical, digital, and personal spheres, it also raises complex ethical, legal, and security questions. Prudent governance should be put in place to manage the transition as AI is introduced into society.

Many use cases already illustrate AI's potential to address social issues. AI systems can help scientists predict environmental changes; for example, Cornell University is using this capability to predict habitat changes and protect certain species of birds.¹⁸ AI also has wide applicability in health care. The Dutch government is using it to identify the most effective treatments for certain patient populations, and cut down on medical mistakes

¹⁴ Vinod Khosla, "Technology will replace 80% of what doctors do," *Fortune*, December 4, 2012.

¹⁵ Executive Office of the President of the United States of America, *Artificial intelligence, automation and the economy*, December 2016.

¹⁶ "Top 20 university majors with highest female ratio," *Sina News*, June 3, 2015, available at <http://edu.sina.com.cn/gaokao/2015-06-03/0952471042.shtml>

¹⁷ *MasterCard Index of Women's Advancement 2016*.

¹⁸ Erin Biba, "Three ways artificial intelligence is helping to save the world," *Ensia*, April 26, 2016.

through analysis of digitalized health records.¹⁹ In the United States, the Las Vegas Health Department is using this technology for public health surveillance, using social media tracking to pinpoint the origins of disease outbreaks.²⁰ AI systems can enhance the safety and efficiency of public transportation and traffic systems. Evidence already shows that AI-enabled autonomous cars can reduce traffic injuries.²¹ Alibaba has collaborated with the government of Hangzhou to make city transportation smarter with AI-directed traffic lights, reducing congestion and speeding traffic by 11 percent in specific areas of the city.²² AI is also being used to predict energy demand and manage energy usage. Early use cases, such as Google lowering energy usage in its massive data centers²³ and the British government managing surges in demand in its grid system,²⁴ point to the possibility of billions in savings for companies and consumers alike.

For all of this potential, it is also an awesome responsibility to manage machines that can learn and make decisions without human direction. These unprecedented capabilities raise many ethical and legal questions that call for serious consideration. Asimov's famous laws of robotics were the first attempt to set out some basic guidelines for robots interacting with humans.²⁵ But the ethical issues brought out by the advent of AI are much more nuanced, and the potential impact is wider.

First, in a world of ubiquitous sensors and AI systems, companies are constantly collecting data about individuals—not only as they use digital devices but as they move through public and personal spaces. In some settings, such as hospitals, this personal information is highly sensitive. This raises questions about who owns personal data, how it may be shared, and how it should be protected from the increasing risk of cybersecurity breaches.

Second, AI may unintentionally discriminate in its decision making. Since the “real world” can be racist, sexist, and biased in many ways, real-world data that is fed into algorithms can also have these features—and when machine learning algorithms learn from biased training data, they internalize the biases. A top AI company experienced such an incident firsthand in 2016 when its experimental “chatbot” offended many Internet users with racist and sexist slurs after the program was trained in online forums. If biased AI is deployed for critical decision-making, its conclusions may even lead to unfair treatment of certain groups of people.

In addition to these ethical concerns, AI's introduction to society also has many legal implications. When accidents and even crimes happen due to AI decision making, who is held liable? Who holds the intellectual property rights to work created by an AI system? With the awesome power of AI, what regulations should be put in place? What legal rights and obligations do AI developers have? These and many other questions will need to be thoroughly debated to create an adequate legal and ethical framework.

¹⁹ “Artificial intelligence will redesign health care,” *The Medical Futurist*, available at <http://medicafuturist.com/artificial-intelligence-will-redesign-healthcare/>.

²⁰ Hope Rees, “Awarding-winning AI app developed at the University of Rochester tracks foodborne illnesses and has implications for public health departments,” *TechRepublic*, March 9, 2016.

²¹ Peter Stone et al., “Artificial intelligence and life in 2030,” *One hundred year study on artificial intelligence: Report of the 2015–2016 study panel*, Stanford University, September 2016.

²² “Aliyun AI ET sets to govern the transportation of Hangzhou,” *Sina News*, October 13, 2016.

²³ James Vincent, “Google uses DeepMind AI to cut data center energy bills,” *The Verge*, July 21, 2016.

²⁴ Valentin Robu, “Why artificial intelligence could be key to future-proofing the grid,” *Robohub*, available at <http://robohub.org/why-artificial-intelligence-could-be-key-to-future-proofing-the-grid/>.

²⁵ Science fiction writer and professor Isaac Asimov devised these rules in 1942. A simplified version is: 1) robots should not harm humans, 2) robots should obey humans unless there is conflict with the first law, and 3) robots should protect their own existence unless there is conflict with the first or second law.

GEOPOLITICAL IMPLICATIONS

The development of AI has been a truly global effort. Further advancement will require international cooperation that promotes broader access to data, algorithms, capital, and talent. But as the global economy grows more digital, many areas of global governance remain a vacuum. The many ethical and security questions associated with the introduction of autonomous systems with greater-than-human intelligence will need to be addressed not only at the national level but through international cooperation.

Furthermore, just as AI-driven automation may create a two-tiered labor market within individual economies, it may enlarge the global “digital divide,” with poorer and less technologically advanced nations falling even further behind on the development curve. Some countries that expect rapid population growth and have been counting on a labor-intensive economic development model may even face new waves of social unrest, as large segments of the population lose their jobs to machines.

Lastly, computer simulation tools are already widely used in war games, and AI will further improve the accuracy and capability of such simulations. But the potential for weaponizing AI is a strong concern. A report commissioned by the US Navy argued that as military robots become more complex, greater attention should be paid to the implications of their ability to make autonomous decisions.²⁶ Stephen Hawking, Elon Musk, and more than 1,000 AI and robotics researchers have signed a letter suggesting a ban on AI warfare, warning of the potential for rampant destruction at the hands of “autonomous weaponry.”²⁷ AI systems, like nuclear energy and nuclear weapons before them, may require strong international agreements to ensure their peaceful use and maintain global security.

CHINA'S PATH FORWARD

It will take a well-crafted strategy to turn today's innovation into a sustainable growth engine for China over the longer term. The government can facilitate this by laying a solid foundation and providing aspirational goals for AI development, spurring private-sector innovation and adoption. Some of the required building blocks include a set of robust industrial, economic, educational, and international policies and frameworks.

INDUSTRIAL AND ECONOMIC FRAMEWORK

Although AI's development is still in the early stages, the technology appears unlikely to follow a linear growth trajectory. The possibility of rapid adoption makes it urgent to ensure that sound industrial policies are put in place. Otherwise China runs the risk of skewed incentives, overinvestment, and oversupply, all of which can destroy value. Although the market will drive the development of AI technology and its applications, the right policy framework can establish a healthy environment for growth.

Strategic priority 1: Build a robust data ecosystem

An abundance of data is a critical ingredient for training AI systems, attracting talent, and accelerating innovation. To build a more robust data ecosystem, China can move to set and implement data standards, open public-sector data for private exploration, and encourage international exchange of data streams.

Standardization is an important precursor to the kind of widespread data sharing and interoperability of systems that will increase the value of the Internet of Things and AI technologies. China is uniquely positioned to lead this effort given the extraordinary amount

²⁶ Patrick Lin, George Bekey, and Keith Abney, *Autonomous military robotics: Risk, ethics, and design*, Ethics and Emerging Sciences Group, California Polytechnic University, prepared for US Department of the Navy, December 2008.

²⁷ Lucas Matney, “Hawking, Musk warn of ‘virtually inevitable’ AI arms race,” TechCrunch, July 27, 2015.

of data potentially available across the nation—and it will need to take the lead to ensure that Chinese-language standards are developed.

For data in specific industries, the government can call on existing regulatory bodies to set out the necessary rules. In the United States, for example, the Securities and Exchange Commission mandated in 2009 that all public companies must disclose their financial statements in XBRL (extensible business reporting language) format, thereby ensuring that public data is machine readable.

To improve the diversity of available data to support AI development, the government can open many more public data sets and take the lead in establishing some industry-specific data sets. In addition to spurring the development of the AI industry, these moves can yield secondary benefits as they improve the quality of public services and reveal new policy insights. The municipal government of New York City, for example, launched its own open data portal, giving citizens access to data on economic development, health, recreation, public services, and more. New York also enacted an open data law in 2012, requiring the government to work with machine readable data and establishing an API (application programming interface) that enables software developers to connect directly to the government system and collect data.

Last but not least, the Chinese government will need to consider the value of international data flows. MGI research has found that cross-border data flows contributed \$2.8 trillion to the global economy in 2014, exerting a larger impact on growth than trade in goods. Further, it finds that both inflows and outflows matter, as they expose an economy to ideas, research, technologies, talent, and best practices from around the world.²⁸ Data is the currency of the future. In medical research, for example, it will be impossible to realize AI's potential without tapping into vast clinical data sets from around the world. Excessive barriers may handicap Chinese AI companies, hampering their ability to develop competitive products in the international markets.

Strategic priority 2: Broaden adoption of AI within traditional industries

Realizing this economic potential in China depends on the actual adoption of AI systems—not just among the technology giants but across China's traditional industries. A great deal of value can be unlocked by improving productivity in these massive sectors. But China needs to address several key barriers.

The first hurdle to overcome involves changing perceptions and creating a sense of urgency about the need to change the way business operates. According to a McKinsey survey, AI is not yet a strategic priority for more than 40 percent of companies in traditional industries in China.²⁹ Consequently, many of them are not yet capturing the data they need to support future AI adoption. Agriculture firms, for example, seldom consider recording detailed information about issues such as planting schedules or the impact of weather on output—but this is precisely the kind of information AI systems could use to unearth valuable insights and efficiencies. By comparison, the United Kingdom, the United States, and Japan have implemented nationwide information systems to capture such data and apply advanced analytics to modern agriculture management.

The second major barrier is a shortage of technical know-how. As discussed above, China will need to focus on developing more elite data scientists, particularly in areas of AI where shortages are becoming apparent. But the talent needed to translate AI knowledge into real-world use cases with real value for business is also in short supply. More business leaders and middle managers will need to have technology skills and the ability to understand and

²⁸ *Digital globalization: The new era of global flows*, McKinsey Global Institute, March 2016.

²⁹ *The rise of the machines: How Chinese executives think about developments in artificial intelligence*, McKinsey & Company, December 2016.

apply data. Like Intel, one Chinese chip maker recognized that the data generated from the manufacturing and testing processes can significantly improve operations and reduce defects. But because it lacked employees with both semiconductor and AI knowledge, the Chinese firm was unable to implement a strategy based on this data.

Last but not least, AI adoption is affected by costs. Purchasing AI systems and hiring the scarce and specialized talent needed to make the most of them is not always cost-effective for Chinese companies. Using technology to streamline manual processes is less urgent when labor costs are low.

AI's greatest economic potential in China is the opportunity to revolutionize traditional industries. If the government sets the stage for this transformation by overcoming some initial barriers to adoption, the market can drive growth.

To facilitate adoption, policy makers should focus on helping the market overcome the three key barriers discussed earlier in this paper: lack of strategic awareness, the costs of adoption, and the shortage of technical know-how.

Some of these issues can be addressed through the traditional economic tools of tax credits and subsidies. The government can also consider pioneering its own adoption of AI systems throughout bureaucracies. This can have a strong follow-on effect that jumpstarts the market, supports government suppliers, and eventually lowers the cost of adoption by accumulating technical experience and talent.

In addition, encouraging adoption of the Internet of Things (IoT) in traditional industries will set the stage for capturing more value from AI adoption, since the IoT connects networks of sensors and devices that can feed AI systems vast quantities of real-world, real-time data. The government can focus on establishing a few successful IoT stories in key economic sectors, alongside its "Internet Plus" policy initiatives, to create models that other traditional industries can follow.

EDUCATION FRAMEWORK

Talent is critical to the development and adoption of artificial intelligence. A robust talent pyramid should have top scientists pushing the boundaries of fundamental AI technology, many developers with the capabilities to create AI applications for real-world contexts, and a large base of workers who will be capable of working alongside AI systems on a day-to-day basis in a wide variety of work settings.

Strategic priority 3: Strengthen the pipeline of specialized AI talent

To address China's talent gap in AI, the government needs to invest in AI-related education and research programs, reorient the education system for a greater focus on innovation and digital skills, and devise an immigration policy to attract the best global talent.

To develop a larger pool of the elite computer scientists needed to advance this technology, the government can invest to establish AI programs and fund AI research labs at top universities. This could involve building AI centers of excellence at top Chinese universities, or sponsoring innovation and research centers to facilitate collaboration among universities, research institutions, and private companies. South Korea's government recently took a firm step in this direction by investing 1 trillion won (\$863 million) to build a state-level, public-private AI research center jointly with leading Korean conglomerates.³⁰ The Canadian government made a similar move by investing more than \$200 million in AI research programs at three Montreal universities.³¹

³⁰ Mark Zastrow, "South Korea trumpets \$860-million AI fund after AlphaGo 'shock,'" *Nature*, March 18, 2016.

³¹ Karen Seidman, "Montreal universities land historic \$213M investment for computer and brain research," *Montreal Gazette*, September 6, 2016.

Many of the experts we interviewed felt that China will need to focus on building a greater culture of innovation to achieve AI breakthroughs. One way to remedy this is to introduce university programs combining AI with other disciplines. Top US universities such as Stanford and MIT have created joint majors combining computer science with humanities subjects with the goal of developing new ways of looking at the world and sparking creativity. These types of programs could inspire new types of AI applications in fields across the economy, including health care, law, finance, and media, to just name a few.

Investing in university programs has long-term benefits, since talent is an important magnet for attracting international companies. Increasingly, large AI developers look to draw talent from academia. Some two-thirds of Google DeepMind's research hires have come directly from academic institutions such as University College London, the University of Oxford, and the University of Montreal.³² Top companies in the field will naturally gravitate toward cities with a large pool of AI talent. For example, both Google and Microsoft have responded to Montreal's growing profile in the field by announcing new investments of their own in local university AI labs as well as local office expansions.³³

In addition to developing more homegrown talent, China will need to work with the best data scientists from around the world and take part in global collaborations. This involves both actively recruiting international experts to work in China and encouraging Chinese AI developers to travel abroad and absorb the latest global research. This may require the government to relax some residential and immigration regulations as well as providing incentives and support.

Strategic priority 4: Ensure that education and training systems are prepared to develop technology skills and retrain large segments of the workforce

Although it may take decades for AI to be widely adopted throughout the economy and society, China will need to prepare for rapid disruptions at the industry level. Certain jobs can disappear in just a few years after a key technology breakthrough. Typists, telephone operators, and darkroom film developers all went largely extinct as technologies made them obsolete.

Helping the labor force in heavily affected industries adapt and gain more relevant new skills will be an ongoing challenge that is critical to maintaining public welfare and social stability. The government will need to proactively identify the jobs that are most likely to be automated and ensure that retraining programs are made available to the segments of the labor force whose livelihoods are at risk. These efforts could involve collaborating closely with vocational training schools and providing educational vouchers to workers.

China will also need to focus on building relevant skills in the future workforce over the longer term. This includes not only building the pipeline of future data scientists and engineers but ensuring that more of the workforce is capable of working alongside technology in a variety of business and professional settings. It will be essential to emphasize science, technology, engineering, and math in schools; even basic education and vocational programs will need to impart data literacy.

Since AI's automation of many routine tasks has the potential to widen the digital divide, it will be critical for the government to monitor the effects on inequality. One aspect of this will be providing equal access to education opportunities. This includes ensuring that female students and students from rural and inland regions receive sufficient exposure to STEM and AI-related courses.

³² Elizabeth Gibney, "AI talent grab sparks excitement and concern," *Nature News*, April 26, 2016.

³³ Rohit Singh Jain, "Microsoft pumps \$7 million into Montreal academia, to expand AI R&D office," *The Tech Portal*, January 19, 2017, available at <https://thetechportal.com/2017/01/19/microsoft-pump-7mn-montreal/>.

SOCIAL AND INTERNATIONAL FRAMEWORK

The advent of AI has the potential to change society in profound ways. Consensus on some of the most pressing ethical and legal questions must be achieved not only domestically but also internationally.

Strategic priority 5: Establishing an ethical and legal consensus among Chinese citizens and in the global community

On the domestic front, preparing the public and achieving a consensus will require a transparent and wide-ranging consultation process. A few legal areas, such as protection of privacy and liability of autonomous vehicles, are of particular importance for AI's development and adoption. The Chinese legislature will need to provide a framework to clear away the legal uncertainties.

Once the legal framework is set, the government will need to establish a regulatory body to monitor and regulate AI activities. Since AI will be adopted in a wide range of industries, this will involve consultation across multiple agencies to draw on their areas of expertise. In health care, for example, the consequences of ill-considered adoption of AI technologies could be severe; the National Health and Family Planning Commission will need to have a strong voice in the development of guidelines.

Internationally, China can take the lead in forming a governing body to promote the peaceful, inclusive, and sustainable development of AI technology. The goal of this international agency should be regulating AI, establishing standards, and developing a code of ethics.

Beyond regulation, China can also take the lead from an economic development standpoint. To ensure that the global digital divide does not become a permanent barrier to prosperity, China can share its AI technology and governance expertise with disadvantaged countries, thereby making an AI One Belt, One Road for a new era.



Artificial intelligence has the potential to fundamentally shape our society for many decades to come. It is a uniquely powerful tool that China can deploy to boost its productivity and maintain its growth trajectory. Moreover, China has the capability and opportunity to lead international collaboration in the development and governance of AI, ensuring that this breakthrough technology will positively contribute to the general welfare of all humanity.

ACKNOWLEDGMENTS

We would like to thank the Development Research Center of the State Council and the China Development Research Foundation for giving McKinsey & Company the opportunity to contribute to the discussion of artificial intelligence and its implications for China. For this paper, we conducted interviews with a broad range of experts from government agencies, academic institutions, and private companies in China and abroad. We also received expertise and insights from colleagues in the McKinsey Global Institute and the Firm's Advanced Industries, Business Technology Office, and Public Sector practices.

We would like to acknowledge with much appreciation the crucial role of colleagues at McKinsey for their time and efforts: Ann Li, Eric Lamarre, James Manyika, Michael Chui, Christopher Thomas, Tong Chen, Lisa Renaud, Lin Lin, Shanshan Shi, Glenn Leibowitz, Claire Colberg, and Mei Ye.

We would also like to express appreciation to the following scholars and leaders in the public, private, and social sectors for their deep insights: Ajay Agrawal, Pam Attinger, Guo Yao, Jui Shang-Ling, Lu Mai, Ma Wenyan, Alex Peng, Shen Zhixun, Wang Xue, Wu Gansha, Simon Yang, Yang Bin, and Yu Jiantuo.



McKinsey Global Institute
April 2017
Copyright © McKinsey & Company
www.mckinsey.com/mgi

 @McKinsey_MGI
 McKinseyGlobalInstitute