

Seizing The Moment For Artificial Intelligence

Andrew Batson
abatson@gavekal.com

In the escalating trade dispute between the US and China, technology has increasingly become the key issue, overwhelming more traditional economic topics like tariffs, deficits and currency valuations. Both countries see their economic future as depending on their success in high technology, and each is worried they will lose out to the other. One of the most intense areas of focus is artificial intelligence, where recent rapid breakthroughs have captivated investors and the media—and where China has emerged as the main US rival.

In this piece, I will try to provide some nontechnical answers to the questions of the moment: What is artificial intelligence anyway, and why is it a hot topic? Why does China seem to be doing so well in artificial intelligence? And how should we think about the rivalry between the US and China to develop this technology? In brief, I think China will do well in artificial intelligence, in part because the technology is now in a phase that plays to its strengths. But it does not make sense to think of the US and China being engaged in an artificial-intelligence “race” along the lines of the US-Soviet space race.

Machines don’t think, but can do useful stuff

A common-sense definition of artificial intelligence, what a layman might understand the term to mean on the basis of reading the news and watching television, is probably something along the lines of: machines that can think like human beings. Artificial intelligence in this sense does not exist, and according to most researchers in the field there is no prospect of it coming into existence any time soon. There are enormous philosophical and technical challenges in understanding how human minds work and replicating those processes in software, and most of these challenges have not been solved.

A more precise definition of artificial intelligence, closer to that used in the industry, would be: the development of computer systems that can perform tasks associated with human intelligence, such as understanding speech, playing games, or driving cars. In a way the term is misleading, because what is being mimicked is not human intelligence itself, but the practical results of intelligence being applied in specific contexts. Artificial intelligence in this narrow sense does very much exist, and its progress is now attracting feverish interest from business, venture capitalists, governments and the media.

The turning point seems to have come around 2014-15. Since then, software programs have been able to match or exceed human performance in tasks that previously could not be reliably performed by machines: recognizing faces, transcribing spoken words, playing complex games. One landmark that had particular resonance in China was the 2016 [victory](#) of Google’s AlphaGo software over a master South Korean player of the board game Go (known as *weiqi* in Chinese); AlphaGo subsequently also defeated China’s top-ranked player. While chess programs have been beating human masters

Artificial intelligence in practice does not have anything to do with making computers think like people

Rather, it is about making machines that are good at performing highly specific tasks

for years, Go is much more complex; the number of potential board positions is traditionally [estimated](#) at 10^{172} , more than the [number of atoms](#) in the universe.

How are these feats possible? Most of what is referred to as “artificial intelligence” in the media is a subset of the field known as machine learning, and in particular a subset of machine learning called deep learning. All software works by following clearly specified instructions, known as algorithms, on how to perform a specific task. In machine learning, the algorithms are not fixed in advance, but evolve over time by building data-driven models.

Recent advances are built on machine learning, whose underlying concepts are decades old

Usually the type of software used in machine learning is called a neural network, because its structure is loosely inspired by the connections between neurons in a human brain. The network takes an input signal and repeatedly processes it into something more useful; what makes the learning “deep” is that there are a large number of “layers” that process the signal. The use of these techniques means that an algorithm can improve its performance of a task by repeated exposure to data. They are particularly useful where writing algorithms the traditional way—by specifying all possible details and eventualities in advance—is cumbersome or impossible. The concept of machine learning dates to the 1960s, and much of the original work underlying today’s approach of deep learning dates to the 1980s and 1990s.

More power, more data

Recent progress is happening by applying more computing power and larger amounts of data

The rapid improvement in the results of specific machine-learning applications in recent years is thus not a result of fresh theoretical breakthroughs. Rather, it has happened because advances in computing power have allowed machine-learning algorithms to run much faster, and the increased availability of very large amounts of structured data have given them much more to work with. The lesson has been that lots of processing of lots of data is required for the algorithms to be effective in finding patterns. This in itself is a sign that machine learning is not very much like human learning: humans can learn quickly from small numbers of examples, by building internal mental models. Machine learning by contrast is a massive and repetitive number-crunching exercise of building up statistical regularities.

Machine learning discovers statistical regularities to help solve problems and make predictions

Researchers in the field sometimes describe machine-learning algorithms as being “narrow” and “brittle.” Narrow means that an algorithm trained to solve one problem in one dataset does not develop general competencies that allow it to solve another problem in another dataset; an algorithm has to be trained separately for each problem. The Go-playing algorithm is not also capable of analyzing MRI scans. Brittle means that the algorithm only knows its dataset, and can break down if confronted with real-world situations not well represented in the data it learned from. An often-used example is facial-recognition software that is trained on databases consisting largely of photos of white men, which then fails to accurately recognize faces of black women.

But while it is important to understand that machine learning is not a magic wand, it would also not do to underestimate its potential. Machine learning is essentially a way of building better algorithms. That means it could be applied to almost any process that already uses software—which, in today’s world, is quite a lot—as well as many new processes that could not be

effectively automated before. The most obvious example is self-driving cars, which can already operate in restricted contexts and could be in general use within a decade. Machine learning is already being used to spot patterns that previously required trained human expertise, such as recognizing financial fraud or early-stage cancers. Because of this broad applicability, enthusiasts call machine learning a “general purpose technology” that, like electricity a century ago, can boost productivity across every part of the economy.

Throwing resources at the problem

The key point is that machine learning has now moved from a pure research phase into a practical development phase. According to Oren Etzioni of the Allen Institute for Artificial Intelligence, all of the major recent successes in machine learning have followed the same template: apply machine-learning algorithms to a large set of carefully categorized data to solve problems in which there is a clear definition of success and failure.

All parts of this procedure are quite resource-intensive. Huge amounts of computing power are required to run the algorithms. The algorithms need huge quantities of data to find patterns. That data must also be first carefully structured and labeled so that the algorithms can draw the right conclusions—for instance, labeling pictures of objects to train an image-recognition algorithm—a process that is extremely labor-intensive. The repeated training and refining of the algorithms also requires a lot of labor by highly skilled workers, whose numbers are necessarily limited. But the reason for the excitement over artificial intelligence is that there is a now a sense that the main remaining constraints on progress are these limitations of resources—and such limitations will be solved over time.

According to its many boosters, China has all of the necessary resources to make progress in machine learning. It has large and well-funded technology companies, including publicly traded giants like Tencent, Alibaba and Baidu, but also private companies with multi-billion-dollar valuations like ByteDance, which runs a popular news app with personalized recommendations, and SenseTime, which specializes in image and facial recognition.

Machine learning has left the fundamental research phase and moved into the practical development phase

Building machine-learning applications require lots of resources, which China has in abundance

Chinese companies still lag in R&D, but are catching up

Ranking of top publicly listed global companies by R&D expenditure, 2017

Rank	Company	R&D spending, US\$ billions								% of revenue
		0.0	2.5	5.0	7.5	10.0	12.5	15.0	17.5	
1	Amazon									11.8
2	Alphabet									15.5
3	Intel									21.5
4	Samsung									7.6
5	Volkswagen									5.3
6	Microsoft									14.1
7	Roche									21.9
8	Merck									25.4
9	Apple									4.7
10	Novartis									19.4
56	Alibaba									10.8
75	ZTE									12.6
81	Tencent									7.8

PwC, Gavekal Data/Macrobond

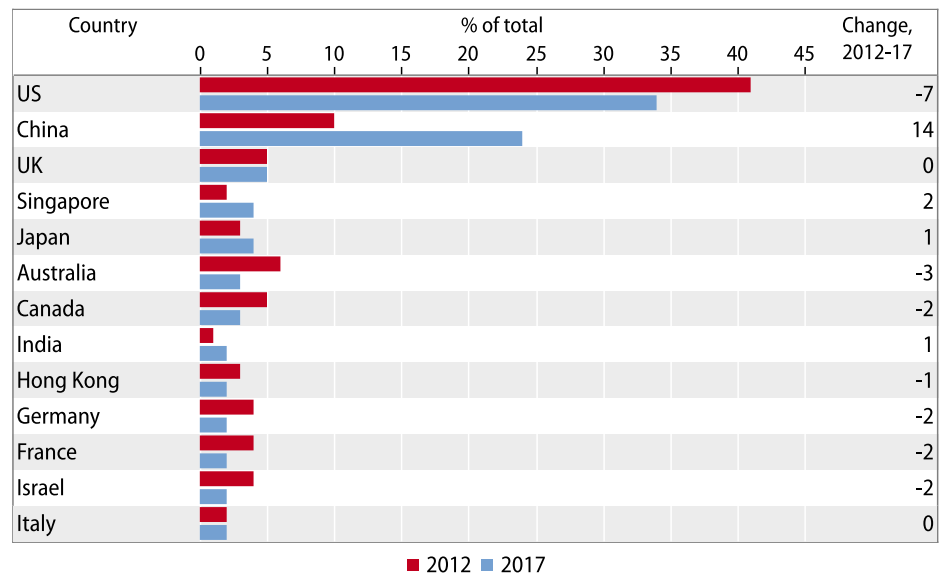
China also has the world's largest population of internet and mobile phone users, who are creating huge amounts of data on daily basis through their interactions with software. It has a huge population of relatively low-cost college graduates, for doing the more repetitive work of categorizing data. And it also has more top artificial-intelligence researchers than any country other than the US; indeed, many of the top Chinese in the field were educated in the US and have worked for US companies.

China also has a government that has decided that artificial intelligence is going to be the key technology of the future, and that will not accept being left behind. An ambitious [national plan](#) released in July 2017 calls for China to lead the world in artificial intelligence theory, technology and applications by 2030 (for detailed analysis of the plan, see these reports by the [Paulson Institute](#) and Oxford University's [Future of Humanity Institute](#)). While it is difficult for government plans to create fundamental research breakthroughs on demand, such plans can be good at mobilizing lots of resources. So to the extent that advances in machine learning are now about mobilizing resources, it is reasonable to think China will indeed be able to make lots of progress.

China's government is mobilizing lots of resources so that it will not be left behind in artificial intelligence

China is playing a bigger role in artificial intelligence research

Country affiliation of authors presenting at AAAI annual conference



Chinese researchers already have a strong and growing presence in the artificial intelligence field

Goldfarb & Trefler (2018), Gavekal Data/Macrobond

China's strength in this area of technology is surprising and shocking to some in the US

The prospect of China being something close to a peer of the US in a major new technology is a shocking development for many Americans. Everyone knows that China's economy has grown rapidly and that it has accomplished a lot. But most of its past successes in technology involve deploying things developed elsewhere, such as mobile phones, wind turbines or high-speed trains, on a large scale. China has a per-capita GDP of roughly US\$8-9,000 at market exchange rates, lower than Mexico or Turkey—and no one is talking about their dominance in the technologies of the future. It is tempting to try to resolve this paradox by focusing on China's state support for artificial intelligence, implying that its advantages are unfair. The rivalry is also not a purely economic one, since there are military uses for machine learning.

The paradox is more apparent than real. China is such a huge, diverse and unequal country that averages are not a good guide to the location of the cutting edge. The reality, as anyone who has visited Beijing, Shanghai, or Shenzhen in recent years can attest, is that the income, skills and education levels of its best people can be comparable to those in the US. That elite of course is not representative of all the hundreds of millions of their compatriots, but neither is the Silicon Valley elite representative of middle-class Americans.

Many of China's top artificial intelligence researchers were educated and employed in the US

The fact that China now has the capability to contribute to cutting-edge research is also in large part a result of its integration with the US: it is the decades of sending top Chinese students to top US universities that have built up the necessary human capital. Rather than say there is a competition between the artificial intelligence sectors in the US and China, it might be more accurate to say that there is a single, global field of machine-learning research that has a significant presence in both North America (Canada also has some top people) and China.

There is no AI race

More fundamentally, it is wrong to think of China and the US as being in a “race” for supremacy in artificial intelligence. Evoking the “space race” with the Soviet Union in the 1960s is the wrong analogy. The space race was about achieving clear technical landmarks defined in advance: first satellite in orbit, first human in orbit, first human on the moon, etc. Today, it's not clear what the technical landmarks for an artificial intelligence race might be. There is a vague goal of “general purpose artificial intelligence,” which means the kind of thinking, talking computers that are familiar from decades of portrayal in science fiction. But there is no race to make one, since no one knows how.

Rather, there are multiple related efforts going on to make progress on diverse sets of specific technical challenges and applications. If China is the first to achieve some technical breakthrough, that does not prevent the US from also doing so, nor does it guarantee that a Chinese company will control the market for applying that breakthrough. Recall that machine-learning applications can be “narrow” and “brittle”: software that is excellent at predicting, say, the video-watching habits of Chinese will not necessarily also dominate the American market. What we can say is that there are economies of scale and scope in machine-learning research: teams of experts who have successfully developed one machine-learning application themselves learn things that will make them better at developing other machine-learning applications (see [this recent paper](#) by Avi Goldfarb and Daniel Treffer for more).

Artificial intelligence is not a prize to be won, or even a single technology

Artificial intelligence is not a prize to be won, or even a single technology. Machine learning is a technique for solving problems and making software. At the moment, it is far from clear what the most commercially important use of machine learning will be. In a way, it is a solution in search of problems. China is making a big push in this area not because it knows what artificial intelligence will be able to do and wants to get there first, but because it does not know, and wants to make sure it does not lose out on the potential benefits. China's development plan for artificial intelligence is mostly a [laundry list](#) of buzzwords and hoped-for technical breakthroughs.

The fact that machine learning is now in a resource-intensive phase does play to China's strengths. There is an enormous amount of venture-capital money and government largesse flowing toward anything labeled "artificial intelligence" in China, and Chinese companies have had some notable successes in attracting high-profile figures in the field to join them. But fears that China will somehow monopolize the resources needed to make progress in machine learning are fanciful. After all, most of the key resources are human beings, who have minds of their own. And many of the key tools and concepts for creating machine-learning applications are in the public domain.

Will today's advantages endure?

The current resource-intensive phase of machine learning may not last forever

It is also not certain that the current resource-intensive phase of machine learning will last forever. As is usually the case when limited resources constrain development, people are trying to find ways to use fewer resources: in this case, refining machine learning so that it does not require so much human effort in categorizing data and fine-tuning algorithms. Some of the current buzzwords in the field are "unsupervised learning," where the machine-learning algorithm is trained on raw data that is not classified or labeled, and "transfer learning," where an algorithm that has already been trained on one dataset is repurposed onto another dataset, which requires much less data the second time around. Progress in these areas could lessen the advantages of China's "big push" approach, though of course Chinese researchers would also benefit from them.

China has a great opportunity to become becoming one of the global clusters of machine-learning research

China's current strength in machine learning is the result of a convergence between its own capabilities and the needs of the technology; since both are evolving, this convergence may not be a permanent one. But China's government is correct to see the current moment as a great opportunity. China was already becoming one of the global clusters of machine-learning research even before the government decided to throw lots of subsidies at the technology. The self-reinforcing dynamics of clusters mean that today's successes will make it easier for China to attract more machine-learning experts and companies in the future.

But political strains are putting pressure on US-China collaboration in artificial intelligence

The biggest loser from this trend, however, is not the US, which already has well-established clusters of machine-learning research, but smaller nations who would also like to become home to such clusters. European countries, for instance, seem to be [struggling](#) to hold their own. The perception that there is a rivalry or race between the US and China ultimately derives from the fact that the two countries are rivals rather than friends. Artificial intelligence may indeed be the first example of a major cutting-edge technology whose development is led by geopolitical competitors—the US and China—rather than a group of friendly nations. The rising tensions between the US and China pose the question of whether a global artificial-intelligence field structured in this way is sustainable, or will be forced to split into national communities. The loss of those exchanges would slow progress in both countries.