

Review of *AI Superpowers: China, Silicon Valley and the New World Order* by Kai-Fu Lee
by Preston McAfee¹

Let me start by concluding that *AI Superpowers* is a great read and well worth the modest investment in time required to read it. Lee has a lively, breezy style, in spite of discussing challenging technology applied to a wide-ranging, thought-provoking subject matter, and he brings great insight and expertise to one of the more important topics of this century. Unlike most of the policy books that are much discussed but rarely read, this book deserves to be read. I will devote my review to critiquing the book, but nonetheless I strongly recommend it.

Lee is one of the world's top artificial intelligence (AI) researchers and a major force in technology startups in China. He was legendary in Microsoft because of his role in Microsoft's leadership in speech recognition. He created Google's research office in China, and then started a major incubator that he still runs. Lee embodies a rare combination of formidable technical prowess and managerial success, so it is not surprising that he has important things to say about AI.

The book provides a relatively limited number of arguments, which I'll summarize here:

- AI will transform business in the near future,
- China's "copycat phase" was a necessary step to develop technical and innovative capacity,
- Success in the Chinese market requires unique customization, which is why western companies were outcompeted by locals,
- Local Chinese companies experienced "gladiatorial" competition, with ruthless copying, and often outright fraud, but the companies that emerged were much stronger for it,
- The Chinese government helped promote internet companies with subsidies and encouragement but not favoritism as widely believed,
- Chinese companies "went heavy" by vertically integrating, e.g. delivering food or bikes, rather than the lean (software only) startup approach of Silicon Valley. The data generated by vertical integration will give Chinese companies a data advantage over their US counterparts,
- Scientific breakthroughs require elite or maverick geniuses, but AI will require an army of implementors, which favors China over the US,
- AI success is mostly about data – everyone has access to the same algorithms – and China has a data advantage over the US. China and the US have a data advantage over

¹ Disclosures: McAfee served as chief economist of Yahoo (2007-12) and Microsoft (2014-2018) and as a research director at Google (2012-2014). McAfee has a contract to advise the FCC on repurposing satellite communications spectrum. In addition, he advises Telescent, a network switching company; Prysm, a blockchain governance company; Merlin, an online employment market; CG, a digital security company in stealth mode; OpenX, a digital advertising exchange; and the Luohan Academy, a not-for-profit research institute created by Alibaba. He also serves on the visiting committee of the MIT Institute for Data, Systems and Society and on the boards of the Pardee-Rand Graduate School and the Mathematical Sciences Research Institute.

the rest of the world. Seven companies dominate AI: Google, Facebook, Amazon, Microsoft, Baidu, Alibaba and Tencent (WeChat).

- AI will roll out in four waves: internet, business, perception and autonomy. We are in the thick of the first two with a start on the third,
- AI will lead to much greater unemployment than most economists expect, displacing perhaps 50% of workers, and having more impact on white collar than blue collar jobs, primarily because robots are difficult and expensive,
- China has more advantages in leading AI than disadvantages, when compared to the US,
- Lee had a personal epiphany due to cancer. This is an interesting story that seems out of place with the rest of the book, almost a Hollywood movie moment,
- Coexistence with AI is going to involve a lot of different patches and workarounds, including education, retraining, transfer payments, and taxes on technology, and also a lot of learning by doing. Lee offers a very thoughtful perspective.

Lee defends these propositions primarily with anecdotes, though he has extensive experience concerning a few of them.

Most people in the developed world regularly encounter AI in three situations: recommendations and ranking of content (e.g. Google or Facebook), digital advertising, and customer service chat. It is not an accident that these are the most popular and successful AI services. All three have a large volume of data to draw on, and more importantly, all have a modest downside risk. When internet search returns bad results, usually the downside is a few wasted moments looking at irrelevant links. When digital advertising shows irrelevant ads – and that is *mostly* what it does – some advertiser’s money is wasted. The downside of failure with chat is usually even lower, because when the chat service can’t help the customer, it rolls over to a human. AI chat services are so good that customers are often unaware they are texting with a bot. All three services regularly fail to work perfectly in spite of voluminous data and billions of dollars in investment. That is acceptable, because the downside risk is low. AI is also extraordinarily good at face recognition, at classifying the content of photographs or video, and is pretty good at language translation. But no AI is widely deployed in a situation with substantial downside risk, like autonomous driving.

Lee takes it for granted that the AI revolution is upon us, indeed discusses Uber without drivers “in a couple of years.” There are several reasons to be skeptical of Lee’s view, which is central to the dire consequences Lee foresees. First, autonomous driving, far from being imminent, seems to be receding. Absent electronic roadway beacons, autonomous vehicles lose their way in snow and indeed, snow builds up on the sensors, rendering them blind. Autonomy also fails frequently in heavy rains. These are technical problems that will eventually be solved, but they have to be solved sufficiently cheaply to replace human drivers. Now, there are already driverless vehicles in contained spaces – in factories, on campuses – and we will surely see many more of those. But the threat of unemployment to general purpose drivers appears quite distant. Moreover, these technical challenges are not the most serious challenges facing AI.

AI makes systematic mistakes. It can be fooled. We shouldn’t be surprised at this; after all, magicians make a living fooling humans and optical illusions fool most of the humans most of

the time. The important point to understand is that the AI methods Lee discusses, deep neural nets, are subject to their own forms of optical illusions.² In some contexts, where the downside is low, the presence of such optical illusions and systematic mistakes doesn't matter much. If an optical illusion causes Netflix to recommend a movie I dislike, I waste a bit of time, and probably Netflix gets better from the experience. But if an optical illusion causes a car to drive into a railroad track, one gets a literal train wreck. These optical illusions would be rare but for vandals, who can exploit the optical illusions in AI to cause damage and mayhem. AI will get better and every vulnerability can be patched. But do the patches lead to an AI no more vulnerable than humans? We don't know for sure, but the best piece of evidence says no. As large email providers, Google, Microsoft and Yahoo each use AI to fight an ongoing battle with vandals (spammers). Every improvement in defenses by email providers is met with more intelligent spam. It appears that highly evolved AI can be tricked for decades. This isn't serious when the downside is low, but is potentially a fatal problem when AI is unleashed on the highways, or in any context where failures are very costly. Far from being a solved problem, AI possesses extraordinary vulnerabilities in any but the most protected contexts. The revolution of AI will roll out slowly in all situations where downsides of failure are substantial.

Lee is keen on Jinri Toutiao, a Chinese company that algorithmically presents news. The US has had such companies. A decade ago I regularly read Techmeme and still scan its headlines daily. Early in its history, Techmeme encountered a problem: its algorithm tended to promote stories involving pictures of scantily-clad women, because many readers clicked on them, and stories could be gamed ("clickbait") to fool the algorithms. In the end, humans were needed to mitigate both problems and are still used today. Algorithmic platforms that aren't curated by humans, like Facebook and Twitter, have their continuing vulnerabilities to manipulation much in the news. Far from being a triumph of AI, news demonstrates the current need for humans and the deep vulnerability of AI to malefactors.

Copycats

Lee argues that unique Chinese circumstances and superior competition by Chinese copycats are the reason for Silicon Valley's lack of success in China, rather than Chinese protectionism. In advancing this argument Lee only compares the US and China. But Silicon Valley companies, particularly Google, Facebook, eBay and Uber, have had great success everywhere *but* China, and Chinese internet companies have had much less success outside of China. If Lee's argument were correct, we would expect to see Chinese internet-based companies flourishing in countries like Indonesia, Iceland or Egypt. To be fair, Huawei and Xiaomi sell phones across the world, but that is likely due to their hardware prowess, and in spite of their software.

Lee strongly defends the copycat history of Chinese innovation. Economists studying development describe a quality ladder – where leading nations sit at the higher rungs of the ladder using the most sophisticated technologies, while less developed nations sit below those rungs, practicing easier and usually labor-intensive technologies. *Every* developing nation has

² See for example <https://spectrum.ieee.org/cars-that-think/transportation/sensors/slight-street-sign-modifications-can-fool-machine-learning-algorithms>.

copied technologies of the leaders, and often had their technologies derided by leading nations. Examples include Japanese cars, watches and cameras, initially ridiculed and ultimately perceived to be commensurate with the highest quality. Lee is correct on this point: Silicon Valley's criticism of Chinese internet copycatting is just a failure to appreciate history.

However, Lee goes further and claims that the competition among copycats, which he repeatedly calls gladiatorial competition, created an advantage of China over the US. There is no doubt that the work environment in China demands long hours and an intense commitment to success, and that Chinese companies have been intensely competitive. However, the kind of competition that Lee documents reveals not a strength of Chinese competition, but a failing. Lee describes, with some implicit admiration, companies that commit fraud and "dirty tricks" to succeed. This is not a recipe for strong companies, but for fraudulent, dishonest ones. The reason we reject the view that "the ends justify the means" is because whatever means are adopted will persist: the means become the ends. Far from producing competitive internet companies, the kinds of competition Lee describes produce Enron, which doesn't bode well. Lee might as well argue that organized crime is good for business.

AI's Effect on Employment

Lee presents a sophisticated but somewhat flawed perspective on the effects of AI on employment and it is a good read. Most commentators on the question of whether AI will lead to major unemployment come from one of two camps, and both camps are naïve. The first says 'technology hasn't caused long-term unemployment in the past, so won't this time.' The second says 'This is the first thinking machine, so this time could well be different, replacing what used to be irreplaceable, at lightning speed.' Lee is closer to the second camp, though more sophisticated in his analysis. Like Lee, both camps are focused on technology as a substitute for human work. However, there are lots of things that people would like to have done – there is no apparent end of demand, and hence no end of potential jobs. Instead, the key question is what the technology asks of the workers. A bulldozer replaces hundreds of workers with shovels, but requires a worker with much higher skills. In contrast, a cotton gin requires a less skilled worker than the workers it replaced. A calculator removes a need for a cashier to perform mathematical calculations, thereby lowering the skill needed. Some technologies are substitutes for human skills, while others, like the bulldozer, are complements, requiring a greater skill level. The key thing to understand about technology is that if it is a complement to skill acquisition, it tends to drive up wages of workers to incentivize the acquisition of skills, while if it substitutes, it drives wages down. That is, the question is not whether AI substitutes for human work – most technologies substitute for work – but whether it tends to increase or decrease the skill level of the workforce.

Is AI a complement or substitute to skill acquisition? It is certainly too early to tell for certain, but Lee himself suggests that the current boom in AI using deep neural nets creates *specific domain* intelligence – recognizing the content of photographs, translating languages, recommending websites – and *not* general intelligence. Indeed, Lee says we may be centuries away from general intelligence (p.143). Lee's description sounds very much like AI is a tool to amplify human expertise, rather than a replacement for human expertise. To put it concretely:

will my taxes be computed by a skilled human with an AI toolkit, or will I send my documents to an AI tax processor? The specialized nature of successful AI programs suggest AI will increase the skill level of the work force, rather than replace skilled technicians with AI and low-skilled workers.

An enormous, perhaps the largest, impact of business AI will be on the organization of employees, simply because humans are the largest component of business costs and one of the least optimized. There are already a variety of AI-based tools rolling out for improving the deployment and effectiveness of workers, including tools to help sales people prioritize which customers to contact and what to discuss; to suggest documents that make people more effective in their jobs; to improve negotiation and pricing with enterprise customers; to aid managers in evaluating their teams; to summarize sets of documents; and to improve project management by reducing the risks of mistakes and wasted resources. Such AI tools tend to be complements to human skills (at a minimum, users need training in using the AI) rather than substitutes.

Finally, Lee broaches but doesn't explore the possibility that AI may revolutionize human education. Our educational systems haven't changed a great deal from the time of Socrates. It is a somewhat astounding fact that the printing press had a fairly minor effect on the operation of higher education. There is certainly the tantalizing possibility that AI systems may create engaging, scalable platforms to supercharge education, along the lines of Neal Stephenson's *The Diamond Age*, where the AI optimizes, personalizes and gamifies the material to amplify every child's natural curiosity. Such systems could improve the capabilities of humanity so dramatically as to render our social and political issues moot.

US vs China

A good portion of the book is devoted to a comparison of the competitive merits of the US and China, specifically on AI but with some broader discussion. Lee treads lightly on this topic; I suspect the book was approved for print by the Chinese Communist Party.³ The rare criticism of the Chinese government is mild. Lee observes that an advantage of the Chinese government is speed. The speed to which he refers is patently evident in China: in a few years, the extensive and inexpensive Beijing subway was built as preparation for the Olympics. This subway covers a Los Angeles-sized area very well, though naturally very crowded at rush hour. Similarly, in half a decade, China connected its major metropolitan areas with bullet trains. The trains cover an area slightly smaller than the continental US and work amazingly well. On the negative side, China has half the world's steel capacity, much of which sits idle, and a handful of empty planned cities. The advantage of command and control systems when speed is critical is probably underestimated by economists, but centuries of history favor democratic, free enterprise economies for the long haul.

US funding of science and universities, amplified by the brain drain, have been enormous US strengths. But as many politicians turned against science and education (e.g. on evolution, on climate change, on stem cells, on religion in schools) and sought to limit immigration, these

³ "China does not intend to use its advantage in the AI as a platform for such colonization" (p.18) strikes me like an addition to please a censor.

strengths are dissipating. In contrast, the Chinese are investing heavily not just in AI, but in space technology, green energy, high speed rail, and medical technology. Relative Chinese investment represents not so much an advantage of the Chinese system of government as a failing of current US politics. At the same time, the world has fully experienced the chilling effects of authoritarian governments on science. Russian astronomers were purged by Stalin,⁴ and China's cultural revolution was devastating to Chinese research.

Lee argues that China's ancient tradition of copying is beneficial. Copying has costs and benefits – it is a benefit that improvements are embodied quickly. But the historical culture of education – that one copies the great master – is surely disadvantageous. In continental Europe, generally an “assistant professor” is someone's assistant, while in the US, assistant professors are expected to perform independent research and be independent thinkers. The US arrangement dramatically increased the number of independent researchers in the US over Europe, and was an important force in the ascendancy of US universities in the twentieth century. China has made great strides in copying the US higher educational system, but the traditions of deference to authority and memorization as learning are impediments to the kinds of independent thinking that Lee envisions. These harmful traditions are surely exacerbated by one party rule.

Lee's discussion of gladiatorial combat by internet copycats inadvertently highlights the absence of due process and the rule of law. Most economists who study development consider rule of law as a critical factor in development. The rule of law is the main bulwark against corruption; the rule of law defends investments and creates the environment for growth. China's weak court system, while permitting advantageous deployments like high speed rail, will likely create headwinds to its future. Will European companies choose American or Chinese AI, when the latter is subject to government whims? Of course, US companies are currently defending against American government attempts to install backdoors and compel access to world data, so the comparison isn't clear cut. But ultimately the rule of law in the US is a substantial advantage to US-headquartered multinationals.

Lee is critical of Silicon Valley's need for a higher purpose and motivation, as opposed to just making money. I think Lee overstates the case two ways – there are many US startups that are focused on money (*cough* Uber), and having social value can be a substantial motivator as well as good public relations.

China's Data Advantage

Lee considers data the greatest advantage that China possesses in AI, and he attributes this advantage to several factors: the size of the population, the greater digital lives of Chinese people, and the prevalence of vertical integration by internet companies. That China's population is large is undeniable, though size doesn't necessarily assure an advantage that extends outside the Chinese market. Indeed, if Lee were right that the Chinese market is so idiosyncratic that only Chinese companies can prosper there, it is likely that the data will be next to useless for services outside the Chinese market.

⁴ https://www.jstor.org/stable/2500602?seq=1#page_scan_tab_contents.

Chinese people, at least those in the urban centers, rarely use cash and perform most transactions with their phones. Lee describes phone-based payments, which are at least a hundred times more extensive in China than in the US, as providing a major advantage in data. China's use of phone-based payments is impressive, in both its scale and the speed with which it was deployed. This amazing deployment happened because so many Chinese could not get credit cards. US credit card companies have a lot of transaction data – Visa transactions are similar in scale to WeChat Pay, and the advantage that China may possess in electronic transactions is locked up in the two major vendors – WeChat Pay and Alipay – limiting the potential value to other companies. Overall, the advantage in financial data of Chinese companies over US companies seems modest or non-existent.

But Lee also considers a data advantage created by internet companies with a foot in the physical world. Lee uses the phrase “going heavy” to mean vertical integration, specifically internet companies investing in real world activities like food delivery. Vertical integration potentially creates a data advantage. But Lee's discussion confounds two distinct phenomena, one of which is immensely valuable, and one that at best has a modest value. Uber or its Chinese competitor Didi create markets: they connect buyers of transport services with sellers of transport services. Similarly, if there are few food delivery options, creating a food delivery market (Lee's example) that connects restaurants and households can be immensely valuable, even if creating the market requires offering transportation. Markets are hard to create, but once created, they are sticky and hard to destroy. Markets convey a durable competitive advantage, for the simple reason that buyers go where the sellers are, and sellers go where the buyers are. Creating markets where there are none is a recipe for the creation of value, one that has nothing to do with data. In contrast, it is usually a terrible idea to vertically integrate and compete against a well-functioning market; instead, it is preferable to focus on one's comparative advantage, and rely on the market for the services it competitively supplies.

Moreover, Lee makes a classic economic error by assuming that what is right for the US is right for China, and vice versa. Labor costs in China are substantially lower than in the US. Using labor to wait in a hospital line or walk dogs may be much more cost-effective in Beijing than in Los Angeles. This fact doesn't vitiate his point about data, but it does blunt his critique of Silicon Valley.

Ultimately, it is not at all clear that China's internet companies will have a data advantage over the US internet companies, or that such a data advantage will be useful outside of China.

How Advantageous is Data?

Lee's perspective is that data advantages will likely persist, because good algorithms and more data tends to beat better algorithms and less data. Lee is certainly right about some AI applications, such as web search. The problem with web search is that what people are seeking – a hot new song, a viral video – changes and the company with the greater flow of data can identify what people are looking for faster. In such changing situations, every provider is starved for data on the new things, and a data advantage is decisive. In other settings, like classifying photographs, understanding human interests, or solving customer service problems, data is less likely to be the constraining factor.

Data depreciates. Some data, like tags in photographs, depreciate very slowly. A photo of the Golden Gate bridge tagged fifty years ago is nearly as useful to AI as one tagged today. Similarly, if you searched for ski resorts last year, you are probably interested in ski resorts this year. But if you searched for pickup trucks last year, probably you are not interested in buying one this year. If you searched for a specific movie last week, you are much less likely to be interested in it this week. And in news, the half-life of data is less than a day. There is a huge difference between tagging photographs, where good algorithms and more data usually beats better algorithms and less data, and advertising, where data depreciates so rapidly that you need a much bigger flow of data to compensate for an inferior algorithm.

It is a delicious irony that Lee, who emphasizes the critical role of data in AI competition, makes a series of forecasts with no data whatsoever behind them.

Conclusion

Lee notes there are seven companies with deep AI capabilities, voluminous data, and extensive resources. Seven is a lot! Many reasonably competitive industries have four or five companies. Moreover, a dozen companies have sufficiently extensive data to build deep profiles of most Americans. As a result, AI technology is likely to be competitively supplied. As AI transforms industries as diverse as appliance manufacture and vacation planning, there will be vigorous competition to supply AI, allowing for innovation, differentiation and customization. That is, the deployment of AI is not likely to reduce competition in other industries, because competing users of the technology have alternative suppliers.

Lee has written an accessible, interesting, and provocative book. Most non-fiction books only manage at most one of those three adjectives. I highly recommend *AI Superpower* to anyone interested in technology, startups, China or the future.