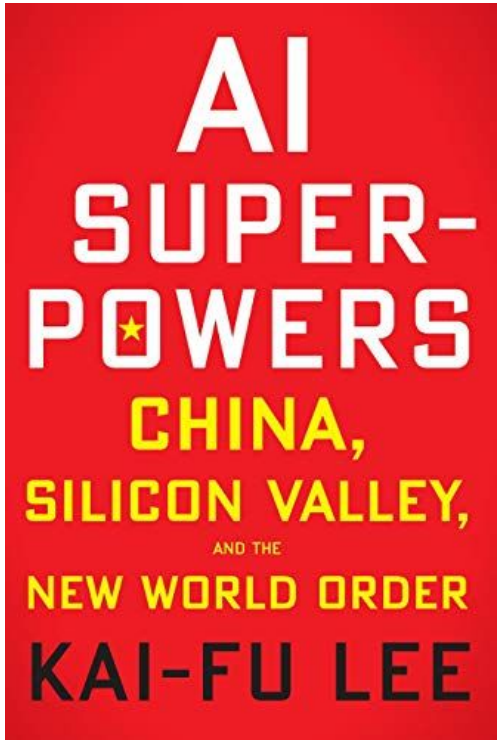



“Worrying about evil AI killer robots today is a little bit like worrying about overpopulation on the planet Mars.”
—Andrew Ng

	<p>Summary: <i>AI Superpowers</i> could be classified as “credible propaganda.” Written by a foremost expert in artificial intelligence, the goal of this book, it seems, is to convince the reader 1) China will win the AI race, and 2) humans need to remember to love one another.</p> <p>Why will China win? China has more people, more data, and a stronger “996” work ethic: “<i>I can tell you that Silicon Valley looks downright sluggish compared to its competitor across the Pacific.</i>”</p> <p>Candidly, this book is heavy on conjecture and light on facts. I suppose that’s to be expected considering Artificial General Intelligence (AGI) doesn’t yet exist.</p> <p>Nevertheless, this book is a good primer on deep learning, artificial intelligence, and whether or not you will lose your job to AI.</p> <p>While Lee’s provocative predictions can, at times, feel completely devoid of facts, there is no doubting the unprecedented wealth creation underway: the seven companies with the largest AI research labs—Google, Facebook, Amazon, Microsoft, Baidu, Alibaba, and Tencent—are also the seven most valuable companies in the world.</p>
--	--

“Instead of simply falling back on a painkiller like a Universal Basic Income [UBI], we must proactively seek and find ways of utilizing AI to double-down on that which separates us from machines: love.”

<p>Kai-Fu Lee</p> 	<p>Author bio: Dr. Kai-Fu Lee is the chairman of Sinovation Ventures, a leading tech investment firm focusing on developing the next generation of Chinese high-tech companies. Before founding Sinovation in 2009, Lee was the president of Google China. Previously he held executive positions at Microsoft, SGI, and Apple.</p> <p>In 1961 Lee was born in Taiwan, later moving to the US at the age of 11. His collegiate work (Columbia, Carnegie Mellon) work focused on speech recognition software, which he brought to the corporate world.</p> <p>Lee is a self-proclaimed workaholic that was diagnosed with lymphoma in 2013 which prompted a reevaluation of his priorities. His deep work in AI (and anti-censorship gadflying) have earned him 50 million followers on Weibo—when he’s not banned. Lee resides in Beijing with his wife and two daughters. I recommend his TED talk, How AI can save humanity, and TED interview.</p>
--	--

“Today, successful AI algorithms need three things: big data, computing power, and the work of strong—but not necessarily elite AI algorithm engineers. . . Data is the core. Having a monopoly on the vest and the brightest just isn’t what it used to be.”

Brief history of deep learning

Lee recaps the history of deep learning, starting in the mid-1950s, by acknowledging some of the “greatest minds in the emerging field of computer science”: [Marvin Minsky](#) (Turing Award!), [John McCarthy](#) (literally coined the term “artificial intelligence”), and [Herbert Simon](#) (Nobel Prize).

In the early 80s, the field of artificial intelligence forked into two camps: **1) rule-based** and **2) neural networks**. The former tried to encode rules based on expert opinions, and the latter: feed algorithms data and let them figure it out.

For the late 1980s and most of the 1990s, Lee claims that AI entered “a prolonged ice age.”

“Deep learning’s big technical break finally arrived in the mid-2000s, when leading researcher [Geoffrey Hinton](#) discovered a new way to efficiently train those new layers in neural networks. The result was like giving steroids to the old neural networks, multiplying their power to perform tasks such as speech and object recognition.”

Turning point: in 2012, a neural network built by Hinton’s team demolished the competition in an international computer vision contest. The improved neural networks had been rebranded as **deep learning** and they hit the mainstream. Google Corp Dev team noticed and acquired both Hinton’s company and DeepMind (\$500m) in 2013.

Lee describes the “core power” of deep learning as: “**recognize a pattern, optimize for a specific outcome, make a decision.**”

China is late to the deep learning party by about 10 years


It is on page 12 that Lee lays out his premise for this book: **China is going to win**. While he acknowledges their late start, Lee argues most Americans still think of China as the copycat laggard. Most of the book is to disavow you of this bias.

This “global shift” of AI-fueled power from US → China is being catalyzed by two transitions, per Lee:

Transition 1: Age of Discovery → Age of Implementation

Transition 2: Age of Expertise → Age of Data

Lee argues that most Americans still think we are in the “Age of Discovery” for deep learning, misled by news headlines of new “breakthroughs” that are simply the application of known technology to *new domains*, with some alpha derived from reinforcement learning and transfer learning (complementary technologies to deep learning).



Deep learning pioneer Andrew Ng has compared AI to Thomas Edison’s harnessing of electricity. Just as 19th century entrepreneurs soon began applying the electricity breakthrough to cooking food, lighting rooms, and powering industrial equipment, today’s AI entrepreneurs are doing the same with deep learning.

“Most of the difficult abstract work of AI research has been done, and it’s now time for entrepreneurs to roll up their sleeves and get down to the dirty work of turning algorithms into sustainable businesses.”

The Age of Data

“Today, successful AI algorithms need three things: big data, computing power, and the work of strong—but not necessarily elite AI algorithm engineers. . . Data is the core. Having a monopoly on the vest and the brightest just isn’t what it used to be.”

While we wait for the next AI step function, the blossoming availability of data will be the driving force behind deep learning’s disruption of countless industries around the world. (this is another reason Lee thinks China will win)

The real crises

What is at stake in all this AI development? Your job.

We’ve heard it before, but Lee’s description is downright haunting, if not a touch hyperbolic:



“Significant as this jockeying between the world’s two superpowers will be, it pales in comparison to the problems of job losses and growing inequality—both domestically and between countries. As deep learning washes over the global economy, it will indeed wipe out billions of jobs up and down the economic ladder: accountants, assembly line workers, warehouse operators, stock analysts, quality control inspectors, truckers, paralegals, and even radiologists, just to name a few.”

Lee goes on to predict that by 2033, AI will be able to replace around 40-50% of the jobs in the United States. As if that weren’t enough, AI often creates a winner-take-all dynamic which will concentrate wealth more and more, while AI-driven automation in factories will undercut the one economic advantage developing countries historically possessed: cheap labor.

The larger threat? Tremendous social disorder and political collapse stemming from widespread unemployment and gaping inequality. Also: all the kittens in the world die. All of them.

China’s internet ecosystem

“Combine the three currents—a cultural acceptance of copying, a scarcity mentality, and the willingness to dive into any promising new industry—and you have the psychological foundation of China’s internet ecosystem.”

To contrast US vs. China business philosophies, Lee shares the story of Jack Ma competing with eBay in the Chinese market (via their acquisition of EachNet) in 2002. While eBay bungled their integration of EachNet with subpar user interface decisions and routing server traffic through the US (costing them speed), Jack Ma was siphoning the marketplace best practices from eBay while catering to Chinese user preferences via Taobao.

Ma understood trust was key so he devised an escrow-esque payment system that delivered payment upon receipt of purchased goods. He also built a messaging platform for increased communication. And then he really stomped on the gas when he launched their freemium model that offered free listings to any Taobao user.

eBay fully retreated from the Chinese market within a few years.

Lee speaks to his own experience launching Google China in 2005 and competing against Baidu. True to form, Google made it hard for Lee to adopt the Google search experience to Chinese users who, they learned, have very different browsing/searching habits from Americans. While Baidu innovated, Google floundered. The reason for this is Silicon Valley companies (and many software companies) don't like to create lots of different versions of their core product. Whenever they "fork" the code base, it makes it harder to manage due to the inherent complexity.

Lee is quick to cite a long list of laudable Silicon Valley firms who have failed, largely due to their own inability to appreciate the nuance of the Chinese market: eBay, Google, Uber, Airbnb, LinkedIn, Amazon.



Reference: Zhou Hongyi, a well-known entrepreneur in China known for his gladiator business mentality and an affinity for heavy artillery. Notably, Hongyi competed fiercely against QQ—the messaging platform of goliath Tencent—and didn't get killed.



Baidu, Alibaba, Tencent (BAT)

The Chinese version of FANG (Facebook, Apple, Netflix, Google) is BAT (Baidu, Alibaba, Tencent). Arbitrary acronyms aside, Lee parades the stories of successful Chinese companies:



It's hard to ignore the unparalleled success of some of China's super apps. Tencent's WeChat, for example, hit 100 million users in just over a year. By year two they had amassed 300 million users. What's unique is that Tencent actually decided to disrupt its own desktop platform, QQ.

As the user base grew, WeChat became a dominant platform for people and companies to build their own apps within WeChat's ecosystem. WeChat Wallet got 5 million bank accounts connected in a day with their infamous "Pearl Harbor attack" on Chinese New Year's Eve.

Market Caps

- Baidu: \$41B
- Alibaba Group: \$692
- Tencent Holdings: \$682B

How do these Chinese companies stack up on the global stage in 2020?

Rank		First quarter	Second quarter
1		Microsoft ▼1,200,000 ^[11]	 Apple Inc. ▲1,576,000 ^[12]
2		Apple Inc. ▼1,113,000 ^[12]	 Microsoft ▲1,551,000 ^[11]
3		Amazon.com ▲970,590 ^[13]	 Amazon.com ▲1,432,590 ^[13]
4		Alphabet Inc. ▼799,180 ^[14]	 Alphabet Inc. ▲979,700 ^[14]
5		Alibaba Group ▼521,740 ^[15]	 Facebook, Inc. ▲675,690 ^[16]
6		Facebook, Inc. ▼475,460 ^[16]	 Tencent ▲620,920 ^[17]
7		Tencent ▲471,660 ^[17]	 Alibaba Group ▲579,740 ^[15]
8		Berkshire Hathaway ▼440,830 ^[18]	 Berkshire Hathaway ▼432,570 ^[18]
9		Visa ▼357,020 ^[19]	 Visa ▲412,710 ^[19]
10		Johnson & Johnson ▼345,700 ^[20]	 Johnson & Johnson ▲370,590 ^[20]



“Of the hundreds of companies pouring resources into AI research, let’s return to the seven that have emerged as the new giants of corporate AI research—Google, Facebook, Amazon, Microsoft, Baidu, Alibaba, and Tencent.”

Side Bar

What are guiding funds: unique to China’s startup ecosystem is government money being invested via “Guiding funds.” How it works is the Chinese government invests in regular venture funds—just as a LP would—but the government caps their own return at, say, 10% and redistributes the other 90% of the gains to the other LPs! This generous ride-along structure has exploded over the last decade: there are now ~1,700 guiding funds in China with a collective war chest of about \$670B (the equivalent of 223 Benchmark Capitals).

Let’s talk about Jack Ma and how the *perception* of entrepreneurs changed drastically in 2014


Lee shares the enduring story of Alibaba founder, Jack Ma. When his company IPO’d in 2014, Alibaba claimed the title of the largest IPO in history (\$21.8B), and Jack Ma became the richest man in China. Jack Ma is known for his quirky, self-deprecating nature (and no coding background) which made him extremely relatable. A legion of entrepreneurs was spawned.

Jack Ma’s current net worth is \$48.2 billion (Aug 2020).



Just five years prior, Lee shares how hard it was to recruit people to work at startups, mostly due to partner and family opposition. To win them over, Lee used to have to wine and dine parents, write letters, and model out financial success for them. After the Alibaba IPO, applicants were (literally) beating down their door. Success sells.

In a book on AI, it's easy to become untethered from practicality. Lee occasionally anchors the reader in some practical applications of the technology, including: 1) autonomous drones, 2) pay-with-your-face systems, and 3) intelligence home appliances. (p.82)



Cornell University

arXiv.org

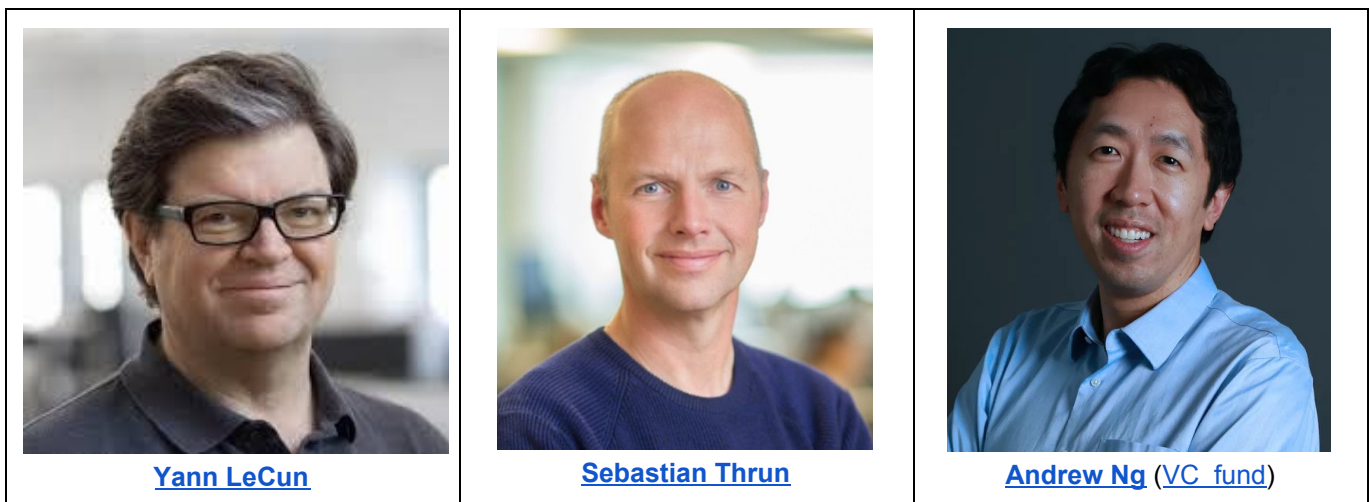
arXiv is a free distribution service and an open-access archive for 1,741,280 scholarly articles in the fields of physics, mathematics, computer science, quantitative biology, quantitative finance, statistics, electrical engineering and systems science, and economics. Materials on this site are not peer-reviewed by arXiv.

Lee references

www.arxiv.org

which he considers a good source of instant publications for scientific papers. It's awesome.

Leading experts in AI: there are many, but Lee calls out 3 in particular from whom Chinese students have eagerly learned from, primarily by devouring translated publications and subtitled lectures.



“Of the top 100 AI researchers and engineers, around half are already working for Google.”

AI Chip Wars & Nvidia

One of the more interesting sections, Lee illuminates the hardware landscape necessary to support AI algorithms. We follow the history of computing from desktop (Intel), to smartphones (Qualcomm/ARM), to AI applications in which Nvidia has crushed it. “These chips are central to everything from facial recognition to self-driving cars, and that has set off a race to build the next-generation AI chip. Google and Microsoft—companies that had long avoided building their own chips—have jumped into the fray, alongside Intel, Qualcomm, and a batch of well-funded Silicon Valley chip startups. Facebook has partnered with Intel to test-drive its first foray into AI-specific chips.

Chinese chip startups like Horizon Robotics, Bittmain, and Cambricon Technologies are flush with investment capital and working on products tailor-made for self-driving cars or other AI use-cases.

→ Lee aptly references the US political headwinds that AI faces, exemplified by the 2017 successful lobbying of Congress to exclude trucks from legislation aimed at speeding up autonomous-vehicle development. **China doesn't have to worry about trucking unions.**

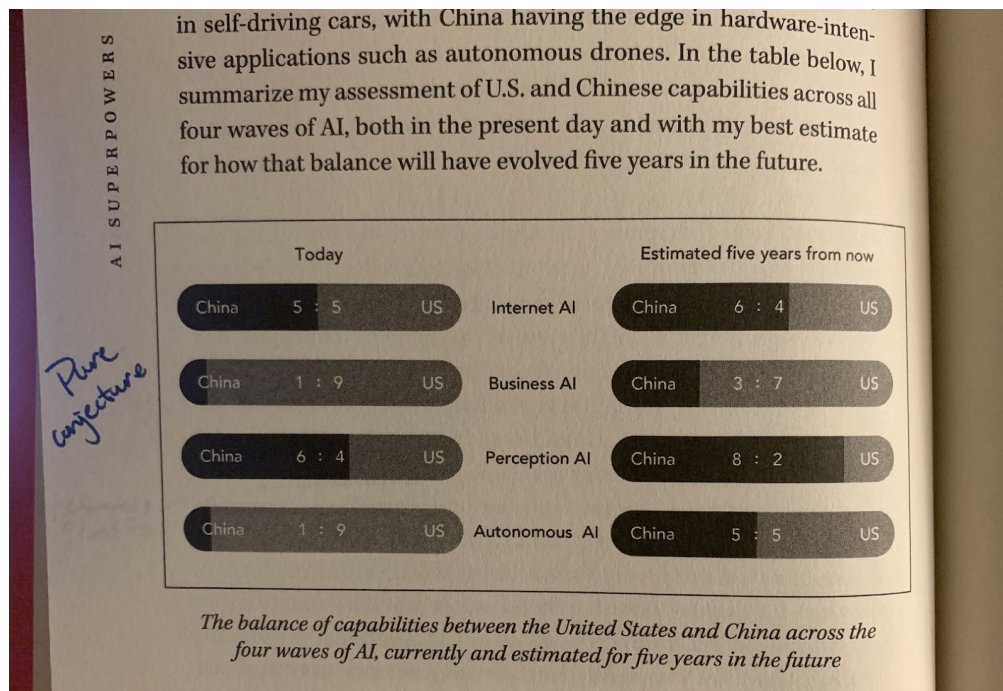
The Four Waves of AI

Lee claims there will be four “waves” of Artificial Intelligence washing over humanity:

1. Internet AI: mostly recommendation engines fueling content consumption
2. Business AI: mining business databases for correlations to predict and optimize stuff
3. Perception AI: converting visual and auditory data in the physical world into value, e.g. Amazon Echo; in China you can pay for KFC with your face via Alipay; simultaneous translation earpieces
4. Autonomous AI: real-world, AI powered robots that can make decisions and shape the world, e.g. drones paint your house, robots pick strawberries, subscription AV

President Trumps speaks Mandarin: iFlyTek—a Chinese company and the most valuable AI speech company in the world according to Lee—was able to create a video of President Trump speaking perfect Chinese by training algorithms on large data samples of his speeches creating a near-perfect digital model of his intonation, pitch and pattern of speech. [Here is the video](#). There is also the use case of translation earpieces that can simultaneous translation which would be **AWESOME**.

Lee gets a little pugnacious on p.136 and starts forecasting the AI scoreboard between the US and China by 2023. **Spoiler alert: China is winning 22-18.** A lot of this book feels like arbitrary c@ck measuring.









Will the future of artificial intelligence be utopian or dystopian?

On p. 141, Lee begins to illuminate the predictions and risks associated with Artificial General Intelligence (AGI). Lots of intelligence people are quotes, many provocative claims are made. But no one really knows, and the expert forecasts—while provocative and headline-grabbing—never seem to reveal their logic or supporting data.

“Worrying about evil AI killer robots today is a little bit like worrying about overpopulation on the planet Mars.”

—Andrew Ng

Utopians (smiley face)			Dystopians (serious face)		
					
Ray Kurzweil	Demis Hassabis	Rodney Brooks	Elon Musk	Stephen Hawking	Nick Bostrom

Ray predicts that by 2029 we will have computer intelligence = human intelligence. By 2045, Ray predicts we will reach the “singularity”—or artificial *superintelligence*. Meanwhile, the challenge that divides the two camps above is known as the “control problem” or “value alignment problem.”

For example, let’s say we challenge artificial superintelligence to **solve global warming**. The algorithm would quickly realize that humans are the primary problem when it comes to global warming. Therefore, remove the humans. That last part is where it gets a bit dicey because the internet is basically *access to everything humans do and know*, and with quantum computing, encryption really isn’t a reliable defense. Depending on the AI’s value function and approach to optimization, it might decide to, say, 1) hack the nuclear launch codes of the country with the *least* technical defense, e.g. Pakistan, North Korea, and 2) fire them at China’s tier 1 cities and US major metros (China and the US produce ~10 and 6 *billion* metric tons of CO2 per year).

Superintelligence might also realize that getting humans to destroy each other would also solve the global warming problem, therefore employing a “instigate World War III” strategy.

Side Bar

Value functions: one of the biggest obstacles to developing Artificial General Intelligence (AGI) is training deep learning algorithms to derive its own value functions from scratch. What’s crazy is a crack team at DeepMind just found [an alternative approach](#) called Learned Policy Gradient (LPG), which is a reinforcement learning (RL) algorithm. **This is a big deal.** LPG was able to create a value function from scratch, meaning it could figure out “what to predict” on its own. Below is a visual of LPG “learning” about its environment.

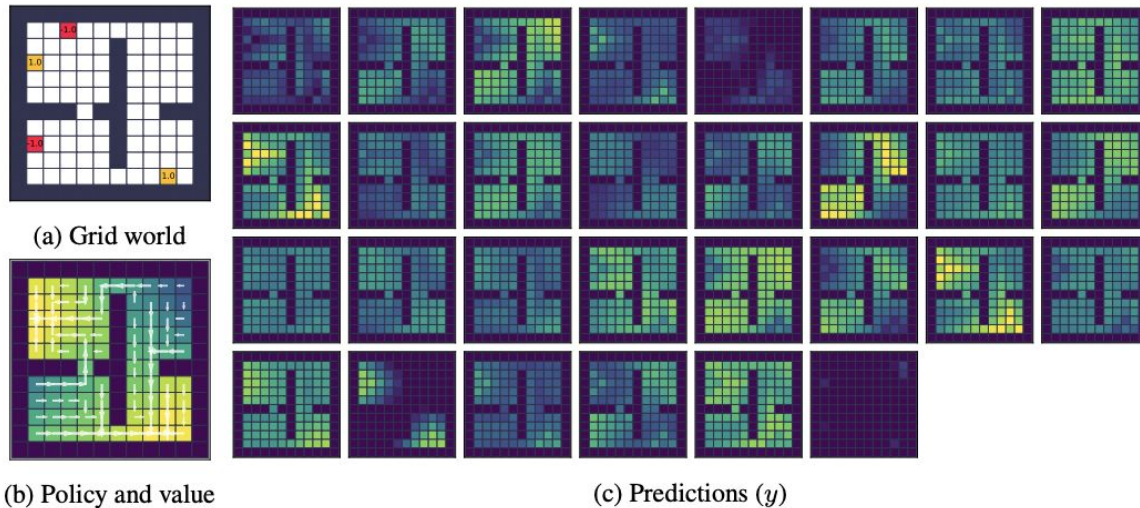


Figure 4: Visualisation of predictions. (a) A grid world with positive goals (yellow) and negative goals (red). (b) A near-optimal policy and its true values. (c) Visualisation of $y \in [0, 1]^{30}$ for the given policy in (b).

Source: <https://arxiv.org/pdf/2007.08794.pdf>

None of this is in Lee's book, but thought it'd be useful to include a RL/AI-related update given the book was written in 2018 and a lot (more) has happened in the last two years.

Provocatively, instead of getting into more details on the facts he references *Folding Beijing*, a dystopian novel by Hao Jingfang that tells a tale of a 3-caste city: 1) First Space: Elite, 2) Second Space: Upper Class, and 3) Third Space, Beijing's economic underclass that toils during the night and sleeps during the day. Third Space does manual, menial tasks that could be automated by robots but the city blesses them with labor so they can earn income. To accomplish this juggling of 75 million people (5m, 20m, and 50m, respectively) the entire city of Beijing distorts and flips upon itself in mechanical inside-out fashion. Third Space works from 10pm to 6am, toiling in the dark (unclear why lighting is not permitted in Third Space). Written by an intelligent female with a PhD in economics and management, many readers absorbed the fear of ruthless economic division due to artificial intelligence.



北京折叠

Is this abhorrent really so different than today? In 2017, the top 1% of Americans possessed twice as much wealth as *the bottom 90% combined*. That's *without* artificial intelligence, so it seems a bit farcical to suggest AI would usher in rampant inequality. It's already here.



Even consulting firms couldn't resist making bombastic predictions. According to PwC, AI will add \$15.7 trillion to the global economy by 2030. 70% of those gains are predicted to go to the US and China.

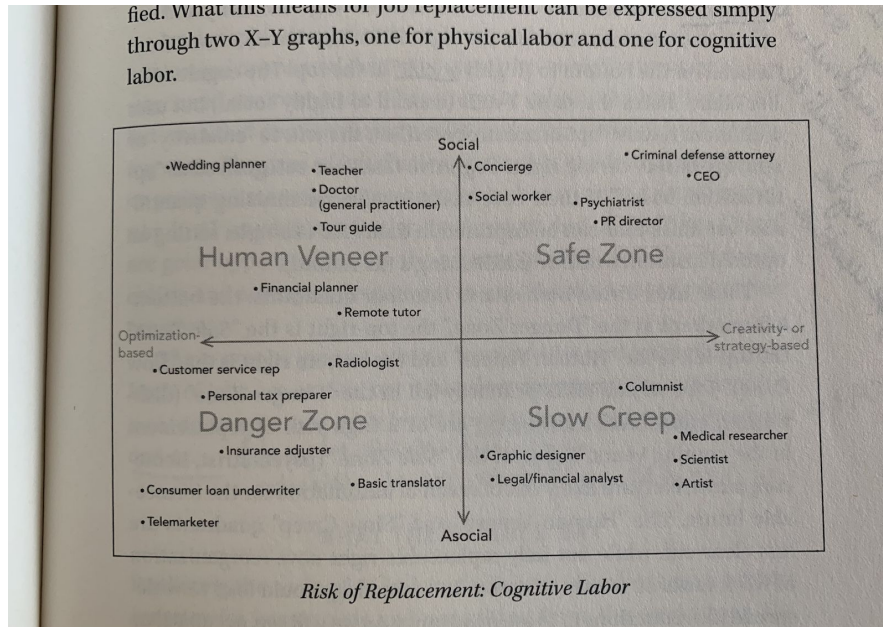
Jobs you don't want your children to have

Lee offers us a 4-box square to help us think about human jobs (cognitive/physical) and the impact of AI:

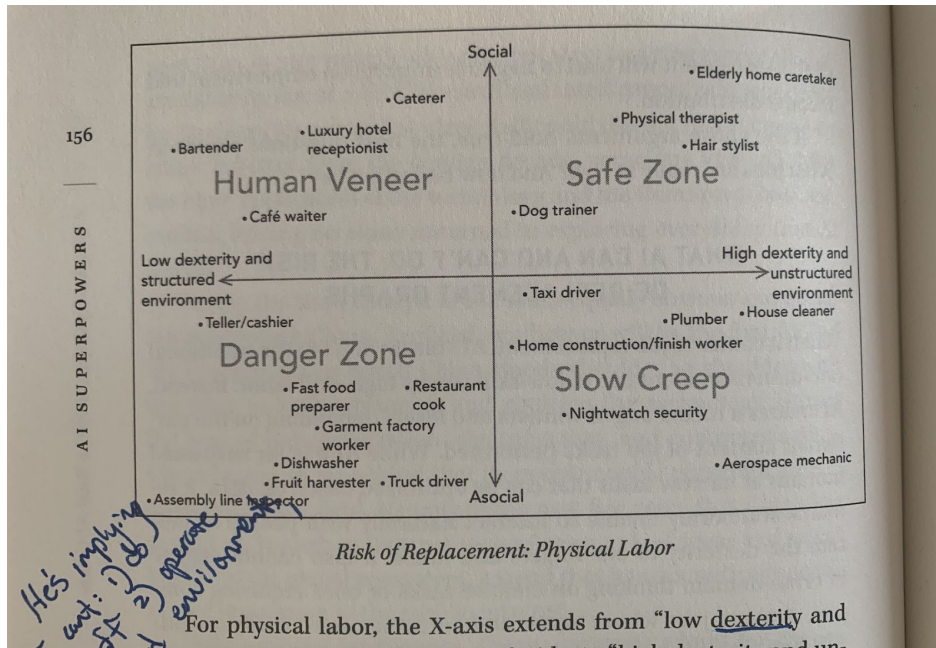
1. Safe Zone: you are okay
2. Human Veneer: AI does the work, you are there to smile and socially shape the experience

DBT VENTURES

3. Slow Creep: your job will eventually be lost
4. Danger Zone: your job is in imminent danger if not eradicated already



In essence, Lee is predicting a future where I can NOT do the following: 1) finger-based dexterity tasks, 2) operate in unstructured environments, 3) socialize, 4) strategize and 5) be creative. So if we all focused on being dexterous artists strategizing on a chaotic commune, we'd be okay.



Overall, we conclude that between 9-47% of jobs could be at risk due to the coming of AI. Of all the stuff humans do **right now**, about **50% are currently automatable**. For example, Lee references the well-known ImageNet computer vision competition where algorithms compete for accuracy in identifying objects. In 2012 the error rate was 16%. In 2017, it was improved to <5% and has since surpassed human performance.



Global management consultant Bain & Co. produced a startling bottom-line conclusion: by 2030, employers will need 20-25% fewer employees. In the US, this would equate to 30-40 million displaced workers.*

*Very similar to impact of COVID-19 in the US, as it were

Dr. Kai-Fu Lee gets personal

In a chapter titled *The Wisdom of Cancer*, Dr. Lee admits that he viewed his life as “an algorithm that optimizes for influence” resulting in his workaholicism. He almost left his pregnant wife Shen-Ling in the delivery room for a work meeting had it not been for an emergency C-section (Lee opens his TED Talk with this story). He was later diagnosed with cancer which catalyzed a drastic reprioritization of his pursuits.

Redistribution of wealth

From p. 200 onward, you can almost feel the conservative blood boiling as Dr. Lee casually surveys solutions to society’s AI-induced unemployment and poverty. The two leading solutions appear to be:

1. **Universal Basic Income (UBI):** everybody gets money from the government each year, say \$10-20k, to invest in their own happiness or re-education, i.e. “VC for the people,” or “personal angel investment”. This idea is gaining support in tech factions like Y-Combinator (Sam Altman et al) and Facebook (Chris Hughes). Most of this solution seems to be aimed at solving the *unemployment* part of the AI double-punch.
2. **Guaranteed Minimum Income (GMI):** almost indistinguishable from welfare for the poor, GMI would establish a human income floor of, say \$26k (the current family poverty line). This solution seems primarily aimed at solving the *poverty* part of the AI double-punch.

Both would be funded by. . . well, it's not very clear. Some say a substantial tax on large companies benefiting from AI, but then I could see companies moving their AI IP into offshore holding companies to sidestep the tax. Others say tax the tech billionaires and then redistribute.

“Instead of simply falling back on a painkiller like a UBI, we must proactively seek and find ways of utilizing AI to double-down on that which separates us from machines: love.”