


BOOK SNAPS™

Zooming In On Your Next Read



AI 2041

By Kai-Fu Lee and Chen Qiufan

Kai-Fu Lee is the CEO of Sinovation Ventures and New York Times bestselling author of *AI Superpowers*. Lee was formerly the president of Google China and a senior executive at Microsoft, SGI, and Apple.

Chen Qiufan (aka Stanley Chan) is an award-winning author, translator, creative producer, and curator. He is the president of the World Chinese Science Fiction Association.

Copyright © 2021 by Kai-Fu Lee and Chen Qiufan • Currency

How Will Artificial Intelligence Change Our World?

AI is a shorthand of which we've all become familiar. It used to conjure up images of robots coming to power over the planet and human beings becoming obsolete at best and enslaved to technology at worst. The truth is that AI is already fairly pervasive in our lives, even though most of us don't understand it or the implications of its proliferation. Nearly four decades ago, coauthor Kai-Fu Lee was applying to Carnegie Mellon University's PhD program and on his application, he described the burgeoning technology in this way:

“Artificial intelligence (AI) is smart software and hardware capable of performing tasks that typically require human intelligence. AI is the elucidation of the human learning process, the quantification of the human thinking process, the explanation of human behavior, and the understanding of what makes intelligence possible. It is mankind's final step in the journey to understanding ourselves, and I hope to take part in this new, but promising science.”

From an earlier than recognized beginning in academia during the mid 1950's, AI saw a very slow evolution in practical application. In the last five years, however, AI has experienced exponential growth in use thanks to a solitary and jaw dropping tipping point. In 2016, engineers at DeepMind built a machine called AlphaGo, designed to play the game of Go, which is “more complex than chess by one million trillion trillion trillion trillion times.” To the astonishment of the (mere) millions of Go fans across the globe, AlphaGo won. Handedly. And it's skyrocketed from there.

“The technology has surpassed humans in speech and object recognition, served up ‘digital humans’ with uncanny realism in both appearance and speech, and earned passing marks on college entrance and medical licensing exams. AI is outperforming judges in fair and consistent sentencing, and radiologists in diagnosing lung cancer, as well as powering drones that will change the future of delivery, agriculture, and warfare. Finally, AI is enabling autonomous vehicles that drive more safely on the highway than humans.” And it's nowhere near finished.



In *AI 2041: Ten Visions for our Future*, coauthors Kai-Fu Lee and Chen Qiufan set out to answer the most asked question about AI: what happens next? They set out to imagine a future world in which AI is fully integrated in a nonthreatening, hopeful, useful way. Every story shared is based on technologies that have a “greater than 80 percent likelihood” of existing by the year 2041 and grounded by the coauthor’s decades of experience in research and development in the AI space. In a fun twist in the tech sector of publishing, the authors shed light on the future of AI through the telling of science fiction stories. Before taking a journey into the future, the authors note that above all, the stories “reinforce our belief in human agency—that we are the masters of our fate, and no technological revolution will ever change that.”

How AI Deep Learning Reaches Goals but Creates Unintended Consequences

The first story centers around an insurance product supported by deep learning AI. A middle class family in Mumbai finds itself in the unexpected position of being able to afford a highly sought after insurance program called Ganesh Insurance. The plan adjusts coverage based on the unique needs of the insured family. In order to access discounts on the plan, the family has to share each family member’s “data link” with the insurer, which gives Ganesh Insurance access to a lifetime of information about each person. The data included “fingerprints, retina signatures, genetic histories, family information, occupations, credit scores, home-buying history, and tax records.” After consent was given, the insurance company was able to “tap into this rich trove of data to personalize its services.”

When the teenage daughter of the family learns that her data is being shared, she is less than pleased but in no time flat, she is freely sharing access to one app or another in order to gain access to the alluring services, including steadily decreasing premiums, being offered in exchange. Trouble strikes when she realizes that the AI system is discouraging her from getting to know a boy in her class that she is interested in knowing better. The boy is from the untouchable class of Indian people, and despite the fact that the caste system was abandoned long before, the boy’s data is like a shadow that follows him, influencing who he can communicate with and where he can go.

The AI is given an objective and it will do “everything possible to achieve that goal.” In the case of the insurance company, it’s objective is to “decrease insurance premiums to the lowest cost possible.” The AI does not take into account anything outside of that end goal, including whether or not the user’s happiness will be adversely impacted. The AI, in general, works spectacularly well and improves the family’s life. They save money, improve their health, and get useful information on a steady stream that makes daily life easier to navigate. But in the case of the teenage daughter, it attempted to derail her chance at love. This was an unintended consequence, or as the author’s call it, an AI externality. In present day life, AI uses deep learning in the same way but as illustrated in the story, it has its limitations.

Deep learning is applied to AI prolifically by massive tech companies like Amazon and Facebook. It drives such data as whether or not a user clicked on something or actually purchased it, which can then be leveraged to impact key business metrics. Financial institutions, including real-life insurance companies, also rely on AI. In that industry there is a veritable race to implement AI as quickly and widely as possible. Using apps for loans, for example, delivers lower default and fraud rates, faster transaction times, and lower costs because there are no humans in the loop.

On the flip side of that coin, there are very real downsides, the impacts of which we don’t yet fully understand. When you click on something, anything, within an AI deep learning system, the supercomputer uses that data to “unconsciously manipulate” you towards advertising that ultimately stands to profit a company. It initiates an addictive cycle which “may narrow your viewpoints, polarize society, distort truth, and negatively affect your happiness, mood, and mental health.”

AI in Healthcare, Including COVID-Inspired Automation Acceleration

The authors wrote *AI 2041* during the COVID-19 pandemic and that experience inspired them to include a story that illuminated how household robotics will reduce the need for human-to-human contact. Further, the authors predict that AI will completely flip conventional medicine on its head with the pandemic being a catalyst for automation in healthcare.

In the fictional story, set in 2041 Shanghai, a woman who is traumatized from previous medical experiences is having an especially difficult time with the mental health implications of COVID-19. To decrease her anxiety she relies on a number of household robots to support daily life. A DeliveryBot, Disinfection Bot, nano super air filtration system, global health tracking tech, communication platforms, and more make life manageable. The systems, all supported by AI of course, made it so easy for her to stay safely inside her home that she planned to never again rejoin society in real life. This is obviously not a great outcome for a member of the most social species on the planet.

In present day life, with the very real pandemic raging on, it has become “clear that AI will reshape healthcare, from speeding discovery of vaccines and drugs to accelerating the integration of technologies like AI diagnostics into existing care.” This comes at an opportune time as much of the healthcare system is becoming digitized, giving AI the massive amount of data needed to be effective. Take medical, wearable devices, and insurance records, for example. As databases get loaded up with more and more data from these sources, AI can be used to improve everything from diagnosis to medical teaching to disease prevention. AI is able to perform at far higher levels than human medical doctors because it can draw from “billions of actual cases” instead of a human clinician’s own limited experience and knowledge.

There are far reaching applications in the field of medical research as well. DNA sequencing is all digitized and perfectly



AI can be used to improve everything from diagnosis to medical teaching to disease prevention.”

suited for AI applications, which will eventually lead to the eradication of many diseases and the development of new drugs and vaccines in the future. Right around the corner lies the potential for “precision medicine” whereby patients are not treated with one-size-fits-all treatment plans. Instead, their unique “medical history, family history, and DNA sequencing” to design treatment plans and select drugs for optimized individual treatment.

Contact tracing, infection control, and more are already being improved with AI integration, spurred into expansion by the COVID-19 pandemic. The authors point to a 2019 study that shows “AI healthcare markets will experience 41.7 percent annual growth to \$13 billion by 2025, in such areas as hospital workflow, wearables, medical imaging and diagnosis, therapy planning, virtual assistants, and most significantly, drug discovery.” The COVID-19 pandemic will very likely accelerate that growth rate even further with diverse positive societal and individual implications. “AI will use big data and individualization data to deliver “precision longevity” by preparing personalized nutrition, supplements, exercise, sleep, medication, and therapy plans for each person.” Experts predict that this could increase average life expectancies by as much as 20 years over the current average.

Autonomous Vehicles and Smart Cities

The state of the vehicle industry has fallen short of our Jetson-inspired flying car fantasies but self-driving landbound vehicles are no longer a figment of our collective imagination. The next story imagines a society in the middle of fully transitioning from human to autonomous driving. In it, we see how both humans and AI have the potential to make mistakes while behind the wheel, but in very different ways.

In 2041 Sri Lanka a young boy is recruited to drive cars in virtual reality. Unbeknownst to him, his driving sessions are actually being used to create data for autonomous vehicles in actual cities in real life. He is taken on a trip to Shenzhen as a reward for his excellent performance and upon arrival is blown away by the level of vehicle automation already in practice. Airports have fleets of automated vehicles operated by a smart control system. Available cars take an optimal path for pick up “based on the passengers’ location and walking speed, to maximize airport efficiency and reduce passengers’ waiting time.” They leave the airport to drive on “smart roads” that have sensors installed along their length that “communicate in real time with the control system of every car and the traffic management infrastructure in the cloud, to ensure safety and orderliness.”

The cars themselves are outfitted with cameras, sensors, and control systems that calculate speed, adjust distance, map optimum routes, and avoid collisions. But there are limitations

that must be supported with human intervention. The most inclusive AI system cannot account for an unexpected natural disaster, for example. The young boy in the story serves as a ghost driver, stepping in where AI may fail, to avert a crisis on the road and supplementing the AI system.

In real life, we are quite a far way from smart roads and fully functioning automated vehicles. Luxury brands offer parking assistance and steering correction using the vehicle’s array of cameras and sensors but the infrastructure is not in place to complement that functionality and take it further afield. Autonomous vehicles are one of the “Holy Grails of artificial intelligence.” The potential for the unexpected and the variety of subtasks and inputs involved in driving make it a function too complicated for current technology to handle safely.

Small steps are being taken and the functionality will mature over time. One major hurdle to overcome is that “AI needs to be trained on large quantities of data that are representative of real driving in many scenarios. However, the number of such scenarios and the degree of variability required are immense—at present there isn’t a feasible way to collect all the permutations of all the objects on the road, moving in all directions, and in all weather conditions.” Nonetheless, the authors predict that vehicles will be approaching full automation by about the year 2040.

The stories, predictions, and current reality of AI in society contained in the pages of AI 2041 paint a fairly rosy picture of the future of humanity. Wealth will be increased broadly, our own human capabilities will be enhanced, and we will live longer, healthier, happier lives with AI than we would have without. Conversely, there are clearly a multitude of complicated challenges and risks that we humans must first rectify before AI permeates every facet of our world.

Humans cannot expect to be “passive spectators in the story of AI.” It will require our very human “creativity, resourcefulness, tenacity, wisdom, courage, compassion, and love” to get right. The very real risks to security, privacy, and job displacements must be balanced, weighed, and mitigated against in the only possible way—our humanity. For their part, the authors are ever optimistic. They fervently believe that “our sense of justice, our capacity to learn, our audacity to dream, and our faith in human agency” will always save the day. We will not simply spectate the story of AI. We will “be the authors of it.” The bottom line to ensure going “boldly where no one has gone before” will be our intentional and calculated symbiosis with AI. We have something that AI will never have—free will. And that is what will ultimately protect and secure the future of humanity.