

# THE AI THREAT

Nouriel Roubini

I HAVE argued that technological progress does not, in the aggregate, destroy jobs. But what happens when that technology is actually intelligent? As science fiction and reality merge in the realm of artificial intelligence, machine learning, robotics, and automation, brace for a cruel twist on the hopes and dreams of inventors ever since the first “mechanical assistants.”


No matter what work you do, artificial intelligence might eventually do it better. Will modern Luddites, for the first time since the original Luddites, finally be correct? The possibility is very real that a tiny top echelon will win while everybody else loses their jobs, their incomes, and their dignity. Mary Shelley’s Frankenstein pales next to this lurking megathreat.

Until very recently the burden of proof stymied believers in the transformative power of artificial intel-

ligence, or AI. A so-called AI winter prevailed during the 1980s and 1990s, as progress was painfully slow and seemed to support skeptics who maintained that computers could never match, much less exceed, the *je ne sais quoi* of human intellectual prowess. Machines improved at doing repetitive things, but deep thinking appeared to remain an exclusively human dimension.

Debate still persists, but the gap between organic and artificial intelligence has decidedly narrowed. Algorithms often ask humans if we are robots before granting access to sensitive websites. By some accounts the gap will soon vanish. More pressure these days compels skeptics to name tasks that computers can *never do*—from bricklaying to neurosurgery. But even with bricklaying, why couldn’t a robot handle it, as there are already AI and 3D technologies printing prefab homes that can build walls far faster than any bricklayer?

**Nouriel Roubini** is Professor of Economics at New York University’s Stern School of Business, Chairman of Roubini Macro Associates ([www.nourielroubini.com](http://www.nourielroubini.com)), and a world-renowned economist. You may follow him on X @Nouriel. This essay is an edited and adapted excerpt from his latest book entitled *MegaThreats: Ten Dangerous Trends That Imperil Our Future, and How to Survive Them* (2022).



Source: Gulliver Image

*The race that humans cannot win...*

Every boost in computing speed and capacity shortens the list. An extreme scenario features the marriage of super intelligent humans with computers that surpass human intelligence, and with robots that have superhuman mechanical abilities. Beyond that point, the world becomes unrecognizable. In effect, we would face a whole new hybrid human species with superior brains and brawn that could displace Homo sapiens, just as we displaced Neanderthal hominids.

If you think your job is safe, think again. Along with desirable boosts in productivity, AI packs unwanted personal and systemic disruption. Before machines get more intelligent than humans, effectively taking control of major portions of technology itself, and technological growth becomes uncontrollable and irreversible—a moment that experts call the singularity—vanishing jobs will strain consumer demand. New jobs may come along to replace them, as in the past, but not if

tailored algorithms can fill those jobs as well. Raising productivity sounds great as the economic pie grows fast until rising inequality and shrinking consumer demand puts more people out of work. When the spiral accelerates, economies hit hard times.

For now, the race is on to deploy artificial intelligence without limit. “This technology will be applied in pretty much every industry out there that has any kind of data—anything from genes to images to language,” AI entrepreneur Richard Socher, the founder of MetaMind, told *The Economist* in 2016. “AI will be everywhere.” Salesforce, a public company that helps other companies reach customers, got the message and acquired MetaMind.

Here is one recent example of technology that reduces costs and eliminates jobs, one that could be directed by AI. Early in 2021, the *New York Post* reported that a 1,407-square-foot gray house with white trim and a front porch on a quarter-acre lot went up for sale in Calverton, New York.

As the first 3D-printed house to obtain approval for sale on Long Island, it made news. Equipment that prints houses and office buildings resembles

a giant hot glue gun on a mechanical arm. Guided by a computer, it dispenses layers of liquid cement in lines to create walls, leaving space for windows and doors. Constructing the frame in nine days required just two workers to monitor the equipment. It cut in half the cost for a conventional house.

*If you think your job is safe, think again. Along with desirable boosts in productivity, AI packs unwanted personal and systemic disruption.*

In July 2021, the Netherlands’ Queen Maxima watched a robot cut a ribbon to open a foot-bridge that spans a canal in the heart of Amsterdam and had been built using a 3D printer.

Touting the bridge’s aesthetic appeal, a spokesman predicted much more to come. “It’s not about making things cheaper and more efficient for us,” Tim Geurtjens said, “it’s about giving architects and designers a new tool—a new very cool tool—in which they can rethink the design of their architecture and their designs.” But now consider the power of AI connected to this scale of 3D printing. When will a computer propose designs without a bridge architect? An architect spends years learning her craft by studying engineering and design. A computer could acquire as much structural knowledge in less than a day.

Do not suppose that creativity requires people. The elusive spark of human ingenuity faces

digital competition. To beat world chess champion Garry Kasparov multiple times in 1997, IBM Deep Blue devised inventive strategies. Yet that was just an opening gambit compared to Deep Mind, a self-teaching algorithm. In 2016, a Deep Mind computer christened AlphaGo mastered a game with more possible moves than there are atoms in the universe. “It studies games that humans have played, it knows the rules and then it comes up with creative moves,” *Wired* Editor-in-Chief Nicholas Thompson told PBS Frontline. In a much-touted contest, AlphaGo outplayed the reigning world Go champion Lee Sedol in four out of five tries.

Game two marked a watershed moment for AI. The 37<sup>th</sup> placement of a piece on the Go board “was a move that humans could not fathom, but yet it ended up being brilliant and woke people up to say, ‘Wow, after thousands of years of playing, we never thought about making a move like that,’” AI scientist Kai-Fu Lee told *Frontline*. Another expert observer suggested, in a sobering coda, that the victory for AI wasn’t so much about a computer beating a human as one form of intelligence beating another. In this battle of brains, neither side enjoys special status.

“You can get into semantics about what does reasoning mean, but clearly the AI system was reasoning at that point,” says *New York Times* journalist Craig Smith, who now hosts the podcast Eye on AI.

A year later, AlphaGo Zero bested AlphaGo by learning the rules of the game and then generating billions of data points in just three days. Deep learning has progressed with mind-bending speed. In 2020, Deep Mind’s AlphaFold2 revolutionized the field of biology by solving “the protein-folding problem” that had stumped medical researchers for five decades.

Besides probing massive volumes of molecular data on protein structures, AlphaFold deployed “transformers,” an innovative neural network that Google Brain scientists unveiled in a 2017 paper. Resolving the protein-folding problem opens the door to significant new bio-medical breakthroughs.

AI-generated artistic initiatives have earned applause. “We have taught a computer to write musical scores,” Gustavo Diaz-Jerez, a software consultant and pianist, told the BBC in 2017. “Now we can produce modern classical music at the touch of a button.” Apart from a rule that playing the music

*Do not suppose that creativity requires people. The elusive spark of human ingenuity faces digital competition.*

cannot require more than five fingers on each of two hands, compositions proceed with very little guidance—and the London Symphony Orchestra has performed several of them. It may only be a matter of time until AI-generated songs make it to the top of the Billboard Hot 100 chart or when an AI-generated novel reaches the *New York Times* bestseller list.

Technology has already channeled Pablo Picasso. A century ago, he painted over an image that was hidden until now. “The nude portrait of a crouching woman has been brought to life by an artificial intelligence-powered software trained to paint like the legendary artist,” NBC News reported in October 2021.

Don’t rule out machines that care. “Xpeng Unveils Smart Robot Pony for Children, Taking It a Step Closer to Its Vision of the Future of Mobility,” the *South China Morning Post* reported in September 2021. “The company said the smart pony, called Little White Dragon, is equipped with power modules, motion control, intelligent navigation and intelligent emotional interaction capabilities.”

*Forbes* book reviewer Calum Chace saw robotic empathy in *A World Without Work* (2020), by Daniel Susskind. “We cannot be confident,” Chace wrote, “that jobs requiring affective capabilities

will always be reserved for humans: machines can already tell if you are happy, surprised, or depressed. Or gay. Some AI systems can tell these things by your facial expressions, and others by how you walk, or dance, or type.” So even nursing jobs for the elderly—previously thought as being available only to humans—may be soon replaced by emotionally intelligent nursing robots.

How many middle-class, white-collar jobs hinge on random access at the right moments to information and skills acquired and stored over the course of a career? The McKinsey Institute, a research arm of the global management consulting firm, concluded in 2016 that compared to the industrial revolution, AI is transforming society “ten times faster and at 300 times the scale, or roughly 3,000 times the impact.”

Researchers Carl Benedikt Frey and Michael Osborne at Oxford University looked into potential job disruption by computers in 702 occupations. Their study, published in 2013, determined that 47 percent of jobs in America were highly vulnerable to substitution by computer capital in the near future. During the Great Depression, 25 percent of Americans were out of work.

Unlike board games that adhere to strict rules, the television game show *Jeopardy!* features puns, slang, red

herrings, vernacular, mischievous wordplay, and obscure associations to elicit factual knowledge on topics from pop culture to the esoteric. No human competitor outperformed erstwhile computer programmer and trivia whiz Ken Jennings, who won *Jeopardy!* 74 times, a legendary streak. Under intense pressure, with the speed of a Google search, he named, for instance, the leader whose brother is believed to be the first known European to have died in the Americas and the disease that prompted U.S. surgeon general Walter Wyman to establish a hospital in Hawaii in 1901 (answers: “Who is Leif Erikson?” and “What is leprosy?”).

Jennings, though, was no match for AI. By his own account in a 2013 TED Talk, IBM’s Watson defeated him handily. He commiserated with Detroit factory workers who became obsolete when robots took their jobs. “I’m not an economist,” Jennings said. “All I know is how it felt to be the guy put out of work and it was freaking demoralizing. It was terrible,” he recalled. “Here’s the one thing that I was ever good at and all it took was IBM pouring tens of millions of dollars and its smartest people and thousands of processors working in parallel and they could do the same thing. They could do it a little bit faster and a little better on national TV and I’m sorry Ken, we don’t need you anymore.” He began to wonder, where

would digital outsourcing of jobs stop? “I felt like a quiz show contestant was now the first job that had become obsolete under this new regime of thinking computers and it hasn’t been the last.”

The philosopher Friedrich Nietzsche envisioned upheaval a century before personal computers arrived. “Every step forward,” he warned in *The Genealogy of Morals* (1887), “is made at the cost of mental and physical pain to someone.”

Robotics and AI firms say you’ll have to wait quite some time before you can own anything remotely similar to Rosey the Robot from *The Jetsons*, the *Washington Post* reported in March 2021. Rosey worked for the cartoon family residing in a future with flying cars and homes elevated to cloud level. “Rosey cooks. She cleans. And she still finds time to play ball with Elroy. Rosey is the ideal maid. Respectful. Even tempered. Does exactly what she’s told. She’s the computer-driven Jill of all trades.” What’s more, Rosey gives sass when suitable. “Beneath the aluminum alloy core beats a battery-powered heart of pure gold.”

We’ll get there. “The biggest problem is safety,” the former chairman of Boston Dynamics, Marc Raibert, told the *Post*. The company has developed agile robots that resemble animals. “The more complicated the robot, the more

safety concerns. If you have a robot in close proximity to a person, and anything that goes wrong, that's a risk to that person," Raibert said.

Decades ago, long before there were any actual robots, the science fiction writer Isaac Asimov proposed three laws to keep us safe from machines we create. Widely quoted since, he first enumerated them in a 1942 short story titled "Runaround." One, a robot may not injure a human being or, through inaction, allow a human being to come to harm. Two, a robot must obey orders given to it by human beings except where such orders would conflict with the First Law. Three, a robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

These laws may not be enough. Writing in the Spring 2016 *Harvard Journal of Law and Technology*, Matthew Scherer weighed inherent conflict when safety competes with completing a task. "Much of the modern scholarship regarding the catastrophic risks associated with AI focuses on systems that seek to maximize a utility function, even when such maximization could pose an existential risk to humanity." In other words, robots may pose a threat by doing what they are supposed to do.

*Robots may pose a threat by doing what they are supposed to do.*

What can go wrong on the job? Plenty. A twenty-two-year-old worker died in a Volkswagen plant in Germany, crushed against a metal plate while setting up a stationary robot in 2015. The same year, a robot arm killed a woman in a Michigan auto plant. A self-driven Uber vehicle

killed a woman in 2018 while its back-up safety driver was streaming an episode of *The Voice*. Authorities let Uber off the hook and charged the back-up driver with negligent homicide.

Killer robots are not the only hitch. People still do some jobs better. Walmart sacked inventory robots in 2020 because "humans can scan products more simply and more efficiently than bulky six-foot-tall machines," according to the *Washington Post*.

An employer can't tell a computer to suck it up and work harder. "Flippy, the burger-flipping robot that threatens to supplant short-order cooks, has taken its first extended break," *USA Today* reported. Flippy, billed as the world's first autonomous robotic kitchen assistant, wasn't to blame. It seems that publicity surrounding her deployment in 2018 in Pasadena, California, created too much demand. Flippy could not keep up. The CaliBurger chain retired Flippy 1.0 and hired more people.

CaliBurger has since deployed Flippy 2.0 in Fort Myers, Florida. There is a hefty appetite for wider use in the fast-food industry where employee turnover can exceed 50 percent a year at a cost of \$3.4 billion in recruiting and training.

Despite some large hurdles, the smart money bets on artificial intelligence. Consumers do not sound surprised. A study by the Pew Research Center in 2017 concluded that three-fourths of Americans find it at least "somewhat realistic" that robots and computers will eventually handle most jobs that people do now.

In Japan, convenience store operator FamilyMart has embraced AI, partly in response to the country's worker shortage. The company intends to open 1,000 fully automated shops by the end of 2024. An unmanned FamilyMart outlet will stock around 3,000 items, the same selection available in shops where people work. A trial store about a third the usual size used 50 cameras to monitor activity and handle payment.

Algorithms are rewriting the art of selling. "Retail Set to Overtake Banking in AI Spending," the *Wall Street Journal* reported in 2021. The website Pinterest assists retailers that use the site to sell their goods. "Everything you can think of in almost every

*Despite some large hurdles, the smart money bets on artificial intelligence.*

part of retail is being powered by AI," Jeremy King, Pinterest's senior vice president of engineering, told the *Journal*. King is also a former executive vice president and chief technology officer at Walmart; he currently sits on the board of Wayfair, which sells furniture and home goods online, and which uses AI to match shoppers with items they might want.

Since 2017, a German e-commerce retailer named Otto has applied AI technology used in particle physics experiments at the CERN laboratory. "It analyses around 3 billion past transactions and 200 variables (such as past sales, searches on Otto's site and weather information) to predict what customers will buy a week before they order," *The Economist* reported.

Home Depot, a brick-and-mortar retailer, taps machine learning to restock shelves. Experts predict that global spending on AI by retailers alone will exceed \$200 billion in 2025, a big jump from \$85 billion in 2021. "You cannot really operate anymore without having the heavy investment in machine learning," Fiona Tan, Wayfair's head of customer and supplier technology, told the *Journal*.

Rudimentary automatons have existed since ancient Greece and Rome ruled the world. The earliest inventors used springs and coils

to make mechanical devices mimic movements by humans or animals.

Meaningful devices that help humans perform tasks proliferated with the industrial revolution in the late eighteenth century. The prospects of machines doing work soon spawned conflict, most famously with the Luddites, who smashed knitting looms. Mill owner William Horsfall paid the ultimate price for automating work. He was shot dead in 1812 while heading home from the Huddersfield town center.

Economist David Ricardo recognized the handwriting on the wall by 1821, thinking seriously about the “influence of machinery on the interests of the different classes of society.” In 1839, Thomas Carlyle (who famously called economics “the dismal science”) fretted about the “demon of mechanism” and its prospects for “oversetting whole multitudes of workmen.” Around the same time, Karl Marx took aim. “Capitalist production,” he warned, “develops technology, and the combining together of various processes into a social whole, only by sapping the original sources of all wealth—the soil and the laborer.”

In 1930, John Maynard Keynes contemplated “Economic Possibilities for Our Grandchildren”:

We are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come—namely, technological unemployment. This means unemployment due to our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour.

Keynes foresaw only “a temporary phase of maladjustment.” He was largely correct, at least until now. “All this means in the long run that mankind is solving its economic problem,” he wrote. “I would predict that the standard of life in progressive countries one hundred years hence will be between four and eight times as high as it is today. There would be nothing surprising in this even in the light of our present knowledge. It would not be foolish to contemplate the possibility of far greater progress still.” He also predicted that technological innovation would lead to a sharp fall in the workweek so that workers could spend most of their time enjoying leisure and artistic and creative activities.

World War II accelerated the pace for automation. Assembly lines built war materiel, newfangled radar tracked aircraft, and researchers at Bletchley Park, England, used advanced mathematics to break secret German naval codes that revealed the

whereabouts of deadly submarines. The brilliant and tragic Alan Turing led the code-breaking initiative. His Enigma machine shortened the war and saved countless lives.

After the war, Turing wrote a paper entitled “Computing Machinery and Intelligence.” Instead of asking whether machines can think, he wondered whether computer responses might seem human by replicating the external manifestations of human thought processes. “This is the premise of Turing’s ‘imitation game,’ where a computer attempts to convince a human interrogator that it is, in fact, human rather than machine,” according to Matthew Scherer in the Spring 2016 *Harvard Journal of Law and Technology*.

Turing imagined a place for artificial intelligence two decades before the term was coined. According to his biographer, Andrew Hodges, “[Turing] supposed it possible to equip the machine with ‘television cameras, microphones, loudspeakers, wheels and handling servo mechanisms as well as some sort of electronic brain.’” Turing proposed, moreover, “that it should ‘roam the countryside’ so that it ‘should have a chance of finding things out for itself.’” We are now not far from satisfying the Turing Test, when a human cannot tell if she is interacting with a machine.

No institution caught on faster than the Pentagon. “New Navy Device Learns by Doing,” the *New York Times* reported in July 1958. “The Navy said the perceptron would be the first non-living mechanism ‘capable of receiving, recognizing and identifying its surroundings without any human training or control.’” In 1962, the first commercial robot took its place on an automotive assembly line. President John F. Kennedy nixed a press conference on the subject of robots and labor and took no action to form a Federal Automation Commission, but he did give a speech about the need to address problems arising from automation.

Anthropomorphic computers got an eerie boost when HAL 9000 commandeered a mission to Jupiter in Stanley Kubrick’s 1968 film, *2001: A Space Odyssey*. Suddenly, humans were dominated by a computer instead of vice versa. HAL’s intentions were suspect. “I know I’ve made some very poor decisions recently,” a deadpan HAL confessed to astronauts aboard the spaceship, “but I can give you my complete assurance that my work will be back to normal. I’ve still got the greatest enthusiasm and confidence in the mission. And I want to help you.” He added a dire warning: “This mission,” HAL announced, “is too important for me to allow you to jeopardize it.” For any artificial intelligence, completing a mission is paramount.

In the years after the film, computers began to alter the nature of work as robots proliferated on shop floors. In 1980, the *New York Times* published an op-ed by Harley Shaiken, a labor activist. Its title: “A Robot Is After Your Job.” Shaiken was blunt: “The introduction of revolutionary new technologies such as robots—versatile computer-controlled mechanical arms—raise two painful possibilities: sizeable losses of jobs and a deteriorated quality of working life.” He advocated an ethos that competed with unfettered capitalism. “The goal, after all, should be a technology that benefits people—not one that destroys them.”

Harvard economist Wassily Leontief amplified a grim message in a 1982 special issue of *Scientific American* magazine. Leontief spelled out issues that have intensified ever since:

There are signs today, however, that past experience cannot serve as a reliable guide for the future of technological change. With the advent of solid-state electronics, machines that have been displacing human muscle from the production of goods are being succeeded by machines that take over the functions of the human nervous system not only in production but in the service industries as well [...] The relation between man and machine is being radically transformed [...] Computers are now taking on the jobs of white-collar workers, performing first

simple and then increasingly complex mental tasks. Human labor from time immemorial played the role of principal factor of production. There are reasons to believe human labor will not retain this status in the future.

Leontief wryly compared humans to horses displaced when the industrial revolution supplied automated horsepower. Artificial intelligence is on track to displace human brainpower in the same way, challenging policy makers to keep up. Yet not until October 2016 did the Obama Administration release a report entitled “Preparing for the Future of Artificial Intelligence.” Both a primer on artificial intelligence and a prescription for interactions between humans and machines, it relied on evidence suggesting that the negative effect of automation would hurt low-wage jobs most.

**T**he artificial intelligence genie is out of the bottle. Its powers are growing, fueled by human nature and free markets. “No matter what monks in their Himalayan caves or philosophers in their ivory towers say, for the capitalist juggernaut, happiness is pleasure. Period,” writes Yuval Harari, the author of *Homo Deus* (2015), a book that posits the marriage of Homo sapiens with artificial intelligence—and super intelligent offspring. By his lights, scientific research and economic activity seek happiness by “producing better pain killers, new ice-cream flavours,

more comfortable mattresses, and more addictive games for our smart phones, so that we will not suffer a single boring moment while waiting for the bus.”

Demographic challenges spur AI to do more. “As China’s working population falls, factories turn to machines to pick up the slack,” the *South China Morning Post* reported in 2021. Don’t look for a person on the shop floor at Midea, a leading maker of home appliances, in Foshan, China. “Human beings have been physically removed from this assembly line, replaced by robots and digital-savvy technicians and engineers operating at a distance.” Once machines get the hang of decisions that remaining people make, those jobs will vanish too.

Efficient competition can bend rules in unsavory ways. “As pricing mechanisms shift to computer pricing algorithms, so too will the types of collusion,” authors Ariel Ezrachi and Maurice Stucke contend in the *University of Illinois Law Review*. “We are shifting from the world where executives expressly collude in smoke-filled hotel rooms to a world where pricing algorithms continually monitor and adjust to each other’s prices and market data.” Surrender scruples or face unpleasant consequences.

**U**neasy lies the head that built the algorithm. Is AI a friend or foe? Will self-learning algorithms replace more human roles, including programmers, than industries of the future can create?

In their book, *The Second Machine Age* (2014), authors Erik Brynjolfsson and Andrew McAfee dismiss the fear that the job market will vanish. They anticipate jobs no one has yet thought of thanks to staggering technological progress. Who foresaw jobs in electronics, data processing, or telecommunications when agricultural and manufacturing jobs began to disappear?

It’s a fair question, but replacing brainpower is different from replacing muscle power. Good jobs that emerged from the decline of manufacturing and rise of services required brains, not brawn. “Knowledge worker” was the category that everyone wanted to join. But now we have lost our monopoly on knowledge. Artificial intelligence can handle desirable jobs better and faster than human brains can handle them. There will be jobs for people, but who will want them?

“The problem is not the number of jobs but the quality and accessibility of those jobs,” says MIT economist David

*The artificial intelligence genie is out of the bottle. Its powers are growing, fueled by human nature and free markets.*

Autor, a prominent expert on the future of work. He reminds a TED audience that automated teller machines (ATMs) slashed the need for bank tellers. The result? Banks built more branches and put would-be tellers to more productive use.

Authors Daniel Susskind and Martin Ford embrace dystopian views in their respective books. They expect AI and robots to fill most jobs. “As we move through the twenty-first century, the demand for the work of human beings is likely to wither away, gradually,” Daniel Susskind warns in *A World Without Work*. Likewise, Martin Ford in *Rise of the Robots* (2015) worries about the threat of a jobless future.

Let’s pause for a moment and look harder at the argument that this time, technological progress will be different. That this is the revolution that will leave us with few and/or worse jobs, unlike all past revolutions. What is different this time?

Industrial revolutions increase productivity. The first revolution introduced steam power. The second revolution launched mass manufacturing. The third revolution harnessed electricity.

The first three industrial revolutions ended many jobs but created more new ones, after some turmoil. None permanently displaced humans. Incomes rose as manufacturing jobs lured superfluous farm workers to move to cities. When manufacturing jobs vanished, the service sector started hiring.

*Replacing brainpower is different from replacing muscle power. Good jobs that emerged from the decline of manufacturing and rise of services required brains, not brawn. “Knowledge worker” was the category that everyone wanted to join. But now we have lost our monopoly on knowledge.*

Today, however, there are fewer places for human workers to go. High-tech firms, the last bastion of fruitful careers, employ far fewer workers than industrial giants in past generations. Facebook—now Meta—is a good example. In late 2021, Meta’s market cap (the combined value of all of its shares) was \$942 billion, making it the world’s seventh most valuable company. But it employed roughly 60,000 workers. Contrast that with Ford Motor Company: its market cap was \$77 billion, but it employed 186,000 workers. Silicon Valley is full of extreme wealth and fast-growing companies, but the tech sector employs far fewer people than older sectors.

And what will happen to Uber drivers and truck drivers worldwide when automobiles drive themselves? Millions of jobs will disappear.

Technology has revolutionized work across the board. Robotic baristas and chefs can displace humans. Recipes are step-by-step instructions on how to cook meals—algorithms, in effect. Express checkout stations replace workers in brick-and-mortar retail stores. Today, e-commerce warehouses rely on robots to move inventory around. Tomorrow, robots and drones will deliver goods to their destinations.

Traditional education capped the typical classroom size at a few dozen students. Nowadays one teacher can reach millions of viewers. Why go to community colleges when top universities come into your home? The experience is not the same, and the outcomes are not identical—as study from home during the COVID-19 crisis showed. But the cost differential is massive, and over time the quality of online education and training will massively improve.

Financial services barely resemble those of a generation ago. Fierce competition has automated tens of thousands of back-office and customer-facing jobs. Computers handle payment services, credit allocation, insurance, capital market support and even asset management. Leading firms advertise

algorithm-based guidance that diversifies and adjusts portfolios faster than humans.

Accounting and legal professionals are looking over their shoulders at electronic job candidates that read and process mountains of documents in seconds. After a pandemic boost, telemedicine has accustomed patients to online health assessments. Computers

*The first three industrial revolutions ended many jobs but created more new ones, after some turmoil. None permanently displaced humans. Today, however, there are fewer places for human workers to go.*

can instantly recall tens of thousands of similar symptoms and diagnoses. Mounting evidence that they discover health problems as reliably as humans moves medicine closer to automating the services of radiologists, nurses, and even physicians. Roles that require human empathy are not exempt. In Japan, hospitals and health care facilities deploy robots to cope with an aging population and a shortage of caregivers.

“If you think being a ‘professional’ makes your job safe, think again,” warned former U.S. labor secretary Robert Reich in an article published by the World Economic Forum. “The two sectors of the economy harboring the most professionals—health care and education—are under increasing pressure to cut costs. And expert machines are poised to take over.”

Researchers Daron Acemoglu at MIT and Pascual Restrepo at Boston University have measured the impact of robotics as it has been introduced in various industries. They found that one additional robot per thousand workers reduces employment by two tenths of one percent, and wages by half a percent. If that sounds trivial, consider the trend. Jobs and incomes are supposed to increase over time. If automation reverses the trend, how do we progress?

MIT's Autor foresees plenty of jobs for highly skilled and very low-skilled workers.

Corporate strategists, neurosurgeons, and health care aides need not make way yet for computers. The vast middle, however, looks problematic. Those jobs “carry out well-defined and codified procedures that increasingly can be done by machines.” Dilberts everywhere, watch out.

Algorithms that learn on their own can do many more jobs once thought exempt from mechanization. Anyone who monitors data, whether doctors, lawyers, teachers, or forest rangers, must compete with mind-boggling computing power that scans and remembers vast amounts of data, and then might propose unconventional responses.

*Traditional education capped the typical classroom size at a few dozen students. Nowadays one teacher can reach millions of viewers. Why go to community colleges when top universities come into your home?*

All this is why the AI revolution may be the first one that destroys overall jobs and wages. Complacency this time—the assumption that once again, the Luddites will be wrong—looks like a fatal mistake. AI encroaches on more jobs than in prior revolutions. It affects jobs across many industries, and it affects knowledge workers just as much as blue collar workers.

Machine learning has accomplished one of the long-term hurdles holding back AI: natural-language processing. By allowing machines to scan vast corpuses of texts and do their own pattern

analyses, AIs have learned how to translate between languages with remarkable success, and how to generate new texts with remarkable authenticity. The subtle grasp of language crosses one of the last obstacles en route to satisfying the Turing Test. “Distinguishing AI-generated text, images and audio from human generated will become extremely difficult,” says Mustafa Suleyman, a cofounder of DeepMind and until recently head of AI policy at Google, as the “transformers” revolution accelerates the power of AI. As a consequence, a large number of white-collar jobs using advanced levels of cognition will become obsolete. Humans won't know that their counterparts are machines.

When I met Demis Hassabis—the other cofounder of DeepMind—he compared the coming singularity to super intelligence that resembles 10,000 Einsteins solving any problem of science, medicine, technology, biology, or knowledge at the same time and in parallel. If that is the future, how can any human compete?

Indeed, AI initially replaced routine jobs. Then it started to replace cognitive jobs that repeat sequences of steps that a machine can master. Now AI is gradually able to perform even creative jobs. So for workers, including those in the creative industries, there is nowhere to hide.

All this is vaulting us even closer to Artificial general intelligence, or AGI, where super intelligent machines leave humans in the dust. Author Ray Kurzweil and other visionaries predict a pivotal moment that will disrupt everything we know. An intelligence explosion will occur when computers develop motivation to learn on their own at warp speed without human direction. There are no limits to how fast or how much they can learn and what new connections they will find. This is what singularity looks like. Human brains

will resemble vacuum tubes in the era of printed circuits, severely limited in capacity.

I asked Demis Hassabis whether ideas once relegated to science fiction look real. He predicts that we are only five major technological innovations and about twenty years away from the singularity.

Unless humans merge with computers, writer Yuval Harari warns, Homo sapiens are finished. They will become obsolete just like Homo erectus, Homo habilis, and other early humans that have long since van-

ished. Enter Homo deus, says Harari, which is smarter, stronger, and immortal so long as knowledge can move from one machine to the next iteration.

Oxford University philosopher Nick Bostrom, the author of *Superintelligence* (2014), ranks artificial intelligence next to giant asteroid strikes and nuclear war as an existential threat to humanity. The late mathematician Stephen Hawking worried that AI “could spell the end of the human race.” That is why he suggested that humans should move to other planets—as the machine will take over not only all jobs but also the human race. Tesla founder Elon Musk

*Machine learning has accomplished one of the long-term hurdles holding back AI: natural-language processing. The subtle grasp of language crosses one of the last obstacles en route to satisfying the Turing Test.*



welcomes AI that controls electric cars his company makes, but putting AI in ultimate charge worries him. “It’s fine if you’ve got Marcus Aurelius as the emperor,” Musk told *The Economist*, “but not so good if you have Caligula.”

No one knows how long it will take for severe structural technological unemployment to make most workers irrelevant. But even the interim looks rocky, prone to negative demand shocks. All signs indicate that AI alternatives will drive down wages and salaries, and that downward drive affects a problem that is already festering.

As people earn less, inequality will grow. Technological innovation is capital intensive, high skill biased and labor saving. If you own the machine or are in the top 5 percent of the human capital distribution, AI will make you richer and more productive. If you are a low- or even medium-skilled blue- or white-collar worker, AI will eventually reduce your wages and make your job obsolete. The trend is already visible in advanced economies where social stability depends on the universal opportunity to achieve success. Data compiled by the Central

Intelligence Agency reveal that income inequality in the United States roughly matches levels in Argentina and Turkey.

Daniel Susskind notes that wealth inequality in the United States is racing out of control. From 1981 to 2017, “the income share of the top 0.1 percent

*An intelligence explosion will occur when computers develop motivation to learn on their own at warp speed without human direction. This is what singularity looks like. Human brains will resemble vacuum tubes in the era of printed circuits, severely limited in capacity.*

increased more than three and a half times from its already disproportionately high level, and the share of the top 0.01 percent rose more than fivefold.” Susskind also cites research into inequality by the scholar Anthony Atkinson, who determined that the top 10 percent of earners saw their wages rise faster than the bottom 10 percent worldwide. Over four decades, Susskind reminds us, CEO

incomes in the United States vaulted from 28 times that of an average worker to more than 376 times in 2000.

Inequality also afflicts the world’s second largest economy. The Chinese government is worried about the growing imbalance between rich and poor. “China’s Media Stars Caught in Harsh Spotlight of Inequality Drive,” *Nikkei Asia* reported in September 2021. “The country’s tech titans have come under the watchful eye of authorities

for practices deemed monopolistic or that run contrary to the common good. Now, even some of China’s most popular stars find themselves in the unforgiving glare of the campaign.”

When the wealthy get wealthier and workers get less, economies suffer from a consumption problem: there isn’t enough of it. Growth eventually may fall as low-income households spend almost everything they have, while the wealthy tend to save more. “As jobs and incomes are relentlessly automated away,” author Martin Ford warns in *Rise of the Robots* (2015), “the bulk of consumers may eventually come to lack the income and purchasing power necessary to drive the demand that is critical to sustained economic growth.”

Although there’s no evidence it actually occurred, the colorful exchange attributed to Ford chairman Henry Ford and United Auto Workers president Walter Reuther helps illustrate the dilemma. The two men were mulling the advent of automation. Ford asked Reuther how robots will pay union dues. Reuther replied, how will Ford

get them to buy his cars? That’s how AI may cause capitalism to eventually self-destruct. A neo-Marxian view of underconsumption spurred by rising inequality that technology exacerbates.

*Technological innovation is capital intensive, high skill biased and labor saving. If you own the machine or are in the top 5 percent of the human capital distribution, AI will make you richer and more productive. If you are a low- or even medium-skilled blue- or white-collar worker, AI will eventually reduce your wages and make your job obsolete.*

This is where the debt burden and AI collide. In a world increasingly driven by AI, the economic pie might become huge for those with highly developed skills that cannot be automated and those who own the means of production.

“Karl Marx was right,” entrepreneur Jerry Kaplan told a tech-savvy audience at Google. “The struggle between capital and labor is a losing proposition for workers. What that means is that

the benefits of automation naturally accrue to those who can invest in the new systems.”

Massive debt disproportionately burdens the people left behind, who live off shrinking paychecks or with public assistance. Less developed countries are highly vulnerable. Those with capital can generate incomes and manage debt. It doesn’t get ahead of them. But for most workers left behind by the rising

machines, a bigger economic pie does not resolve the growing debt problem; it only gets worse.

As a human, I root for people. As an economist I must ask what is the most efficient use of resources? How can we assure the long-term continuation of progress and take care of workers? Priorities conflict.

Over the next decades there will be winners in parts of Europe, China, and North America. Many other countries will become losers, swept under by technological unemployment and drowning in debt they cannot service much less ever repay. Polarization will pit the rich against the poor.

Enter the new precariat, educated and semi-skilled workers who lose careers to AI and end up in gig work with unstable income and no benefits. They will go from job to job with no future, falling through a fraying safety net. Then what happens? As incomes fall, they may try to borrow more. Debt loads increase as income gaps widen. An ugly situation that currently looks intractable gets worse with no sure remedy in sight.

Education geared to a world with increased automation might salvage

some incomes, but a shrinking job market limits potential. Unfortunately, more education is no panacea to the onslaught of AI. Returns on education were higher when a modest upgrade in skills could lead to a better job and more income. When entry-level jobs require advanced degrees, however,

*In a world increasingly driven by AI, the economic pie might become huge for those with highly developed skills that cannot be automated and those who own the means of production.*

upgrades short of that won't change the picture. Not everyone has enough talent and inclination to program computers, explore databases, improve AI, write successful novels, or become entrepreneurs. When AI displaces skilled work, the returns from education become smaller.

If people cannot work, then what? The answer looks like a political minefield: it's time to tax the winners. A tiny contingent will reap the lucrative rewards that AI bestows. Taxing robots as if they were human sounds appealing but really amounts to almost the same thing: taxing the owners of the machines.

If we adjust taxation for this brave new world the next question centers on redistribution that is vital to sustaining demand for the goods that robots produce. One option surfaced during the 2020 presidential campaign in the United States: universal basic income

(UBI) that lets consumers consume. Besides replacing lost income, proposals include more robust public services under the banner of universal basic provision (UBP). Twists abound, including community service in exchange for UBI. We could give each individual a share of ownership of all firms. Then they would receive capital returns even if their labor incomes are challenged. If you think about it, this is a form of socialism where every worker owns the means of production. It is not hard to envision a scenario where people who demonize these choices today as socialist will clamor for them when algorithms perform brain surgery and prepare fast food.

*When AI displaces skilled work, the returns from education become smaller.*

Any of these options will lead to pitched political battles. If we squabble long enough, computers may get to decide how to divide the economic pie. By then, let's hope they have empathy.

"The most important question in twenty-first-century economics," says Yuval Harari, "may well be to do with all the superfluous people. What will conscious humans do once we have highly intelligent non-conscious algorithms that can do almost everything better?" In some dystopian scenarios, "superfluous" people disappear. UBI lets them play video games all day and use drugs that eventu-

ally precipitate "deaths of despair." Drug overdoses caused more than 100,000 deaths in the United States in 2021. Alternatively, young men may become sexually inactive incels who don't reproduce themselves and thus disappear. Our dystopian future may conflate Orwell's Big Brother, Huxley's Brave New World and the Hunger Games.

We are racing toward destiny. Human nature propels us forward. I won't sugarcoat a story about

super intelligent artificial offspring. I do not foresee a happy future where new jobs replace the jobs that automation snatches. This revolution looks terminal. The flowering of artificial intelligence might alter human life beyond recognition.

Earth may be lucky to reach the intelligence explosion of the singularity. Will a deadly pandemic finish us before the transition to machines is complete? Will climate change destroy the planet before rational machines come to the rescue? Will we suffocate under a mountain of debt? Or will the United States and China destroy the world in a military conflict as competition to control the industries of the future becomes extreme? Indeed, who controls AI may become the dominant world superpower. ●