

THE WALL STREET JOURNAL.

This copy is for your personal, non-commercial use only. To order presentation-ready copies for distribution to your colleagues, clients or customers visit <https://www.djreprints.com>.

<https://www.wsj.com/articles/the-worlds-that-ai-might-create-11571018700>

JOURNAL REPORTS: TECHNOLOGY

The Worlds That AI Might Create

Artificial intelligence will have a profound impact—on our jobs, our health and possibly our very existence. But that’s where consensus ends.

By Michael Totty

Oct. 13, 2019 10:05 pm ET

When it comes to artificial intelligence, there’s a clear consensus: It is a growing presence in our offices and homes. But the consensus ends when you ask the next question: What will it mean?

To some experts, an AI world means more jobs, and more interesting ones; to others, it means a devastating loss of employment opportunities. To some, it means a deadly threat to human existence; to others, it means better health and longer—perhaps much longer—lives. To some, it means a time when AI can help us make smarter decisions; to others, it means the destruction of our privacy.

How are experts looking at the same present and arriving at such different and contradictory futures? Here’s a look at five scenarios, and the paths that getting there might take.

Many jobs will disappear, and won’t be replaced

As artificial intelligence becomes more powerful, a lot of current jobs are doomed to disappear.

University of Oxford researchers in 2017 estimated that nearly half of all U.S. jobs were at risk from AI-powered automation. Other forecasts come up with different estimates, but by any measure, the number of lost jobs is potentially huge.

Automation has already made manufacturing, mining, agriculture and many other industries much less labor-intensive. One study estimated that from 1993 to 2007, each industrial robot replaced 3.3 workers. With about 2.5 million robots in industry now and more than three million expected by 2020, the trend is expected to accelerate, and manufacturing could lose up to 20 million jobs by 2030, according to a study this year by Oxford Economics.

While many economists may believe that AI will create more jobs than it destroys, this time history doesn’t serve as a guide. Unlike in the past, when new fields of economic activity arose to provide lots of new jobs, that isn’t happening these days.

Why is this time different? For starters, AI is able to take over almost any routine work, including jobs that might otherwise be created by new economic tasks. And as it becomes more capable, it will increasingly be able to take on many nonroutine ones as well.

What’s more, while previous tech revolutions created jobs for unskilled workers, many or most of the new jobs that will be created by AI will require education and skills that most of those who lose their jobs will lack. It’s possible they can be retrained, but it’s unlikely that former truck drivers will become machine-language programmers.

In addition, government and economic policies are reinforcing the trend. Capital investments in computers and robots get a tax break, while labor is taxed. And the new economy is dominated by innovative, fast-growing companies that succeed with far fewer employees. “A

company like Facebook, its business model doesn't have much need for humans," says Daron Acemoglu, a professor of economics at the Massachusetts Institute of Technology. "What it needs are better algorithms."


Office-support and customer-service jobs often rely on routine, repetitive tasks, and will be among the first to fall to AI, as systems using voice recognition and natural-language processing continue to improve.

Then there are jobs that robots can't take over completely but that have elements that could be easily automated. McKinsey Global Institute estimates that about a third of the tasks in 60% of the occupations it studied fall into this category. Many employers will cut their workforces and take the savings.

The jobs that will be the most difficult to automate are those that require empathy and "people skills." For instance, "college professors can be replaced more easily than kindergarten teachers," says Jamais Cascio, an author and futurist. "Heart surgeons can be replaced more easily than nurses. Clothing designers can be replaced more easily than hairdressers."

In other words, because AI automates many cognitive chores, a college degree and a white collar won't be enough to shield those jobs, either.

"Everyone has the bias that if you have education and skills, you're going to be protected from automation. That in many cases is quite wrong," says Martin Ford, author of "Rise of the Robots: Technology and the Threat of a Jobless Future." "It's not about skill. It's about the nature of the work."



'We need to be asking some very basic, fundamental questions,' says Marina Gorbis, executive director of the Institute for the Future. PHOTO: EOIN HARDY

There will be plenty of jobs (just different ones)

It's true that many jobs will be lost in the AI revolution, just as in previous waves of automation. But history *is* a guide, and once again, even more jobs will be gained.

McKinsey has forecast that the equivalent of 400 million jobs world-wide could be automated by 2030. At the same time, it projected that productivity gains and growing consumer demand would mean as many as 890 million *new* jobs, more than enough to offset the losses.

That is, over the next several decades, AI not only will create the need for many new jobs and new types of jobs, but it also will transform existing jobs in ways that make them easier, safer and more productive. What's more, the increased productivity will make possible both more leisure time and the opportunity for more meaningful and creative work.

“It will change the nature of jobs,” says Peter Schwartz, a senior vice president of strategic planning at Salesforce. “Some will go away, but we’re going to create many more.”

There are all sorts of reasons to expect AI to be responsible for a job boom. For one thing, developing and implementing AI systems creates a growing demand for data scientists, roboticists, machine-learning specialists, cybersecurity experts and other highly skilled workers. Just as nobody could have anticipated that the Industrial Revolution would create millions of new jobs in factories, mills and mines, there will be millions of jobs that we can’t even imagine today that will spring from the AI revolution.

Artificial intelligence also requires a host of new companion jobs to “train, explain and sustain” AI systems, says Jim Wilson, managing director of information technology and business research at Accenture Research. Tens of thousands of workers around the world have full- or part-time jobs training machine-learning algorithms by manually identifying pictures of cats or picking out tumors on radiology images. Those jobs are only a hint of what’s ahead.

Perhaps most important, rather than replace jobs, robots and other AI systems will work alongside humans and enhance their knowledge and skills. Scientists at a bioscience company, Mr. Wilson says, use robotic lab equipment to run experiments more precisely, enabling researchers to conduct in a single year tests that would take them 100 years on their own. Such jobs will still need humans to handle tasks requiring creativity and problem solving, such as designing new experiments, or for manual chores that require quickly adapting to changing situations.

Even in highly automated factories, people and robots working together are more productive than either working alone. There, cooperative robots, or “cobots,” handle heavy lifting or repetitive tasks while their human co-workers take care of duties that require dexterity, on-the-fly problem solving and mobility in unpredictable environments.

AI will also change many jobs beyond recognition. Truck driving, for example, is among those jobs at the greatest risk once AI-powered autonomous vehicles hit the road, perhaps as early as the next decade. But despite what doomsayers fear, jobs driving trucks won’t go away. Even the most capable self-driving truck will have trouble navigating city streets or suburban neighborhoods. For those situations, a driver in a remote control center—much like drones are piloted now—might guide the vehicle in and out of the neighborhood and on to the freeway, where it becomes almost fully autonomous. “The skill set now is Grand Theft Auto,” Mr. Schwartz says.

Even as it transforms many jobs and creates millions more, there’s no question that lots of workers will be displaced in the process. But that doesn’t have to lead to higher unemployment.

McKinsey forecasts that AI will contribute to a 2% increase in productivity over the next decade as goods and services are produced at lower cost. The wealth created by that higher productivity could be used to boost employment and salaries in teaching and child and elder care, which face a growing demand and require a uniquely human touch. It could also go toward expanding investment in infrastructure improvements and in making the economy more sustainable, adding millions of new jobs. “It’s an incredible stimulus,” says James Manyika, chairman of the McKinsey Global Institute.

Higher productivity has another upside that’s almost unimaginable in our workaholic society: more leisure time. There’s no reason three- or four-day workweeks and shorter workdays, with no loss in purchasing power, couldn’t be the norm.

“Why do we need to work five days a week if we could avoid it?” says Yvette Wohn, an assistant professor of informatics at the New Jersey Institute of Technology.

Our very existence is threatened

Here’s the truly nightmare scenario of artificial intelligence: It kills us all.

This isn’t just a movie plot (though it is that, too). To many serious thinkers about AI, this is a

real threat that those developing AI systems need to plan for now.

The danger isn't from robots that will seek to control and destroy humans. No, it's more benign-sounding than that.

"The real risk isn't AI turning evil like in the movies, but turning competent and accomplishing goals that aren't aligned with ours," says Max Tegmark, a professor of physics at the Massachusetts Institute of Technology and a co-founder of the Future of Life Institute, which researches ways to make AI safer.

How might it happen? One possibility is that researchers succeed in creating a humanlike AI system—what is called artificial general intelligence, or AGI—that is capable of learning on its own and that could then design itself to be even more intelligent. In this event, which researchers refer to as the singularity, the machine could improve so rapidly that it turns into a superintelligence that is beyond our ability to monitor or control.

Such a computer would be able to commandeer resources, such as automated factories or the computerized financial system, to achieve its objectives with indifference to the consequences, and regardless of whether its mission matches up with what humans want.

This difficulty in aligning AI and human values is akin to what tripped up King Midas, the Sorcerer's Apprentice and everyone in fairy tales who dealt with genies.

Stuart Russell, a professor of computer science at the University of California, Berkeley, imagines assigning a super AI to quickly come up with a cure for cancer. The system digests all the existing literature on the disease and comes up with millions of possible treatments—all of them untested. To test their effectiveness, the AI introduces cancerous tumors in every person on Earth and begins medical trials, some safe and some not.

The problem, Prof. Russell says, is that it's almost impossible to anticipate every path a super AI might take to achieve its objective. "If you leave anything out, the AI system will find a way to take that thing you left out and shove it to infinity to help optimize the thing that you said you wanted," he says.

Couldn't we just turn off a superintelligent AI before it starts to do harm? It turns out that's not easy to do. Prof. Russell notes that an AI that's hellbent on achieving its objectives would also realize that being shut down would prevent its ability to succeed and would try to stop any effort to pull the plug. (See HAL in "2001: A Space Odyssey.") Instead, he and others warn, it's necessary to build in safeguards long before a humanlike artificial intelligence is created.

"If humanity unleashes superintelligence without careful safety engineering," Prof. Tegmark says, "the default outcome is trouble."

We'll be healthier and live longer—maybe a lot longer

AI is going to become superintelligent and kill us? Not likely, many researchers say. Scientists not only don't know how to create a humanlike AI, they aren't likely to figure it out soon, if ever.

No, the opposite is much more feasible: AI is going to make it possible to live longer, healthier lives. And perhaps a lot, lot longer.

The reason is that instead of becoming our master, artificial intelligence will become our servant. By tapping the power of artificial intelligence to find patterns in enormous amounts of

data—about our health, our genes, our environment and our lifestyles—doctors will be able to make better diagnoses and recommend more effective treatments. Researchers will better understand how diseases work and devise more targeted and personalized ways to treat them. And everyday users will have powerful diagnostic tools that can spot early warning signs of illness.

“There is no area of medicine that will be spared from AI’s impact,” says Eric Topol, executive vice president and a professor of molecular medicine at the Scripps Research Institute in La Jolla, Calif.

Start with the doctor’s office. Physicians, in theory, already have access to previously unimaginable sources of health information: electronic medical records, radiology and lab reports, the patients’ fitness trackers and the results of genetic tests. But by themselves, it’s almost impossible for doctors to draw meaningful insights from all that information.

AI systems will fill that gap.

They already have shown in various studies that they can analyze medical information and come up with a correct diagnosis as well as or better than clinicians. And those diagnostic skills will get immeasurably better as our use of AI systems improves.

Patients themselves will also get medical help from AI-powered “health personal assistants” that will advise—and prod—users to take more healthy actions. For instance, Dr. Topol describes how diabetics could carry a virtual medical coach that takes information from glucose monitors, sleep and activity trackers and other sources and provides guidance on what they should be eating and what activities would help control blood sugar.

Finally, AI will help researchers identify new medical treatments and, perhaps, unlock the secrets of aging.

The body’s decline as it ages is a complex biological and chemical process that involves nearly all its systems, organs and cells. For longevity researchers to understand how these parts interact means crunching an enormous amount of data, and sophisticated AI techniques are increasingly being put to the task.

“Our goal is to have everyone be young for as long as possible,” says Alex Zhavoronkov, chief executive of Hong Kong-based Insilico Medicine, which is using AI to try to solve the problem of aging.

AI will be a constant companion

It won’t be long before AI will be following us everywhere.

The path to a ubiquitous AI isn’t hard to imagine. Artificial intelligence is an all-pervasive, general-purpose technology, more like electricity than, say, the airplane. Like electricity, it eventually will be integrated into all aspects of our lives, homes, cars and offices, though in ways that are far more disruptive and far-reaching.

AI will drive us to work in our autonomous cars, and once we’re there it will manage calendars, screen and interview job candidates, run meetings, and even take on some management tasks such as forming work teams and assigning projects. Back at home, smart devices will react automatically to changing temperatures, noise levels and air quality, change lights and music to fit our mood and help children with their homework.

“At a certain point in the near term, referring to a building as AI-enabled would be as silly as referring to one as electrified today,” says Mr. Cascio.

Some people may find it hard to imagine that they will turn over all these things to AI. But they’ll do it for a simple reason: It will make our lives easier by managing all the scattered details we otherwise would have to pay attention to ourselves. An AI assistant, for instance, would track any delays to your spouse’s arriving flight and, taking account of traffic to the

 JOURNAL REPORT

- [Read more at WSJ.com/Alreport](#)

 MORE IN ARTIFICIAL INTELLIGENCE

- [Creating a Robot-Human Bond](#)
- [AI Learns to Write Fake News](#)
- [Relationship Help From an Algorithm](#)
- [Debate: Should You Be Able to Sell Your Personal Data?](#)

airport, give you an alert to leave in five minutes —after reminding you the night before to charge your electric car. “We’re overwhelmed and looking for something to help us focus our attention in the most fruitful way,” Mr. Cascio says. “It’s not so much laziness as it is triage.”

We’ll also trust that our AI companion will help us make better decisions, and more quickly. Partly that’s because it will have access to far more information than we can have, much as the Waze driving app knows there’s traffic congestion ahead. Even today, few people are likely to ignore their GPS instructions and decide

they know best. That will be more so—and about so many more things—in the future.

Imagine that you’ve just read about the latest unrest in the Middle East or a trade war with China and decide to unload your stocks ahead of a possible financial meltdown. Your assistant, knowing that you’re in the throes of a temporary, irrational panic, would prevent you from executing the sale until you’ve had a chance to calm down.

It’s easy to see this is where we’re heading. The bigger question is, what will it mean? “What kind of life is it, when more of these decisions are taken by the algorithms?” Yuval Noah Harari, an Israeli historian, asks in a TED interview describing this scenario.

One possibility is that turning over decisions and actions to an AI assistant creates a “nanny world” that makes us less and less able to act on our own. It’s what one writer has called the “Jeeves effect” after the P.G. Wodehouse butler character who is so capable that Bertie Wooster, his employer, can get by being completely incompetent.

A simple example most of us can identify with: Using GPS for directions can reduce our ability to find our way around. “I used to pride myself on being able to navigate, but now that’s slipping,” says Patrick Lin, director of the Ethics and Emerging Sciences Group at California Polytechnic State University. “It’s hard to see the benefit of offloading that ability to technology.”

Then there’s the threat to privacy. The more we rely on AI, the more personal information we’re giving to the AI software. “Anything with the word ‘smart’ in it needs data to learn from,” says Azeem Azhar, who advises companies on the impact of AI and who publishes the Exponential View newsletter. “As soon as you have a smart something in your home, you have to start thinking very hard about what happens to your data.”

Finally, in this future, how we interact with the world may very well change as we try to accommodate our behavior to our indispensable, ever-present AI companions. Users describe “barking” commands to the Alexa voice-controlled assistant; we could start barking to the people in our lives as well. Or consider that automated customer-support systems require speaking in a mechanical voice and can be tripped up by accents and other individual quirks.

“We need to be asking some very basic, fundamental questions,” says Marina Gorbis, executive director of the Institute for the Future, a Palo Alto, Calif., research and forecasting organization. “How do we shape these machines to be more like humans rather than making us more machine-like?”

Mr. Totty, a former Journal Report news editor in San Francisco, can be reached at reports@wsj.com.

SHARE YOUR THOUGHTS

Will the growth of artificial intelligence ultimately be a net positive or a negative outcome for humankind? Why? Join the conversation below.

-
- **College Rankings**
 - **College Rankings Highlights**
 - **Energy**
 - **Funds/ETFs**
 - **Health Care**
 - **Leadership**
 - **Retirement**
 - **Small Business**
 - **Technology**
 - **Wealth Management**

Copyright © 2019 Dow Jones & Company, Inc. All Rights Reserved

This copy is for your personal, non-commercial use only. To order presentation-ready copies for distribution to your colleagues, clients or customers visit <https://www.djreprints.com>.