

The New York Review of Books

How Robots & Algorithms Are Taking Over

Sue Halpern

APRIL 2, 2015 ISSUE

The Glass Cage: Automation and Us

by Nicholas Carr

Norton, 276 pp., \$26.95

In September 2013, about a year before Nicholas Carr published *The Glass Cage: Automation and Us*, his chastening meditation on the human future, a pair of Oxford researchers issued a report predicting that nearly half of all jobs in the United States could be lost to machines within the next twenty years. The researchers, Carl Benedikt Frey and Michael Osborne, looked at seven hundred kinds of work and found that of those occupations, among the most susceptible to automation were loan officers, receptionists, paralegals, store clerks, taxi drivers, and security guards. Even computer programmers, the people writing the algorithms that are taking on these tasks, will not be immune. By Frey and Osborne's calculations, there is about a 50 percent chance that programming, too, will be outsourced to machines within the next two decades.

In fact, this is already happening, in part because programmers increasingly rely on “self-correcting” code—that is, code that debugs and rewrites itself*—and in part because they are creating machines that are able to learn on the job. While these machines cannot think, per se, they can process phenomenal amounts of data with ever-increasing speed and use what they have learned to perform such functions as medical diagnosis, navigation, and translation, among many others. Add to these self-repairing robots that are able to negotiate hostile environments like radioactive power plants and collapsed mines and then fix themselves without human intercession when the need arises. The most recent iteration of these robots has been designed by the robots themselves, suggesting that in the future even roboticists may find themselves out of work.

The term for what happens when human workers are replaced by machines was coined by John Maynard Keynes in 1930 in the essay “Economic Possibilities for our Grandchildren.” He called it “technological unemployment.” At the time, Keynes considered technical unemployment a transitory condition, “a temporary phase of maladjustment” brought on by “our discovery of means of economizing the use of labour outrunning the pace at which we can find new uses for labour.” In the United States, for example, the mechanization of the railways around the time Keynes was writing his essay



CCI/Art Archive/Art Resource

Artwork for the cover of a 1959 issue of the French science fiction magazine Galaxie

put nearly half a million people out of work. Similarly, rotary phones were making switchboard operators obsolete, while mechanical harvesters, plows, and combines were replacing traditional farmworkers, just as the first steam-engine tractors had replaced horses and oxen less than a century before. Machine efficiency was becoming so great that President Roosevelt, in 1935, told the nation that the economy might never be able to reabsorb all the workers who were being displaced. The more sanguine *New York Times* editorial board then accused the president of falling prey to the “calamity prophets.”

In retrospect, it certainly looked as if he had. Unemployment, which was at nearly 24 percent in 1932, dropped to less than 5 percent a decade later. This was a pattern that would reassert itself throughout the twentieth century: the economy would tank, automation would be identified as one of the main culprits, commentators would suggest that jobs were not coming back, and then the economy would rebound and with it employment, and all that nervous chatter about machines taking over would fade away.

When the economy faltered in 1958, and then again in 1961, for instance, what was being called the “automation problem” was taken up by Congress, which passed the Manpower Development and Training Act. In his State of the Union Address of 1962, President Kennedy explained that this law was meant “to stop the waste of able-bodied men and women who want to work, but whose only skill has been replaced by a machine, moved with a mill, or shut down with a mine.” Two years later, President Johnson convened a National Commission on Technology, Automation, and Economic Progress to assess the economic effects of automation and technological change. But then a funny thing happened. By the time the commission issued its report in 1966, the economy was approaching full employment. Concern about machines supplanting workers abated. The commission was disbanded.

That fear, though, was dormant, not gone. A *Time* magazine cover from 1980 titled “The Robot Revolution” shows a tentacled automaton strangling human workers. An essay three years later by an MIT economist named Harley Shaiken begins:

As more and more attention is focused on economic recovery, for 11 million people the grim reality is continued unemployment. Against this backdrop the central issue raised by rampant and pervasive technological change is not simply how many people may be displaced in the coming decade but how many who are currently unemployed will never return to the job.

Unemployment, which was approaching 10 percent at the time, then fell by half at decade’s end, and once more the automation problem receded.

Yet there it was again, on the heels of the economic collapse of 2008. An investigation by the Associated Press in 2013 put it this way:

Five years after the start of the Great Recession, the toll is terrifyingly clear: Millions of middle-class jobs have been lost in developed countries the world over.

And the situation is even worse than it appears.

Most of the jobs will never return, and millions more are likely to vanish as well, say experts who study the labor market....

They’re being obliterated by technology.

Year after year, the software that runs computers and an array of other machines and devices becomes more sophisticated and powerful and capable of doing more efficiently tasks that humans have always done. For decades, science fiction warned of a future when we would be architects of our own obsolescence, replaced by our machines; an Associated Press analysis finds that the future has arrived.

Here is what that future—which is to say now—looks like: banking, logistics, surgery, and medical recordkeeping are just a few of the occupations that have already been given over to machines. Manufacturing, which has long been hospitable to mechanization and automation, is becoming more so as the cost of industrial robots drops, especially in relation to the cost of human labor. According to a new study by the Boston Consulting Group, currently the expectation is that machines, which now account for 10 percent of all manufacturing tasks, are likely to perform about 25 percent of them by 2025. (To understand the economics of this transition, one need only consider the American automotive industry, where a human spot welder costs about \$25 an hour and a robotic one costs \$8. The robot is faster and more accurate, too.) The Boston group expects most of the growth in automation to be concentrated in transportation equipment, computer and electronic products, electrical equipment, and machinery.

Meanwhile, algorithms are writing most corporate reports, analyzing intelligence data for the NSA and CIA, reading mammograms, grading tests, and sniffing out plagiarism. Computers fly planes—Nicholas Carr points out that the average airline pilot is now at the helm of an airplane for about three minutes per flight—and they compose music and pick which pop songs should be recorded based on which chord progressions and riffs were hits in the past. Computers pursue drug development—a robot in the UK named Eve may have just found a new compound to treat malaria—and fill pharmacy vials.

Xerox uses computers—not people—to select which applicants to hire for its call centers. The retail giant Amazon “employs” 15,000 warehouse robots to pull items off the shelf and pack boxes. The self-driving car is being road-tested. A number of hotels are staffed by robotic desk clerks and cleaned by robotic chambermaids. Airports are instituting robotic valet parking. Cynthia Breazeal, the director of MIT’s personal robots group, raised \$1 million in six days on the crowd-funding site Indiegogo, and then \$25 million in venture capital funding, to bring Jibo, “the world’s first social robot,” to market.

What is a social robot? In the words of John Markoff of *The New York Times*, “it’s a robot with a little humanity.” It will tell your child bedtime stories, order takeout when you don’t feel like cooking, know you prefer Coke over Pepsi, and snap photos of important life events so you don’t have to step out of the picture. At the other end of the spectrum, machine guns, which automated killing in the nineteenth century, are being supplanted by Lethal Autonomous Robots (LARs) that can operate without human intervention. (By contrast, drones, which fly without an onboard pilot, still require a person at the controls.) All this—and unemployment is now below 6 percent.


Gross unemployment statistics, of course, can be deceptive. They don’t take into account people who have given up looking for work, or people who are underemployed, or those who have had to take pay cuts after losing higher-paying jobs. And they don’t reflect where the jobs are, or what sectors they represent, and which age cohorts are finding employment and which are not. And so while the pattern looks familiar, the worry is that this time around, machines really will undermine the labor force. As former Treasury Secretary Lawrence Summers wrote in *The Wall Street Journal* last July:

The economic challenge of the future will not be *producing* enough. It will be providing enough *good jobs*.... Today...there are more sectors losing jobs than creating jobs. And the general-purpose aspect of software technology means that even the industries and jobs that it creates are not forever.

To be clear, there are physical robots like Jibo and the machines that assemble our cars, and there are virtual robots, which are the algorithms that undergird the computers that perform countless daily tasks, from driving those cars, to Google searches, to online banking. Both are avatars of automation, and both are altering the nature of work, taking on not only repetitive physical jobs, but intellectual and heretofore exclusively human ones as well. And while both are defining features of what has been called “the second machine age,” what really distinguishes this moment is the speed at which technology is changing and changing society with it. If the “calamity prophets” are finally right, and this time the machines really will win out, this is why. It’s not just that computers seem to be infiltrating every aspect of our lives, it’s that they *have* infiltrated them and *are* infiltrating them with breathless rapidity. It’s not just that life *seems* to have sped up, it’s that it has. And that speed, and that infiltration, appear to have a life of their own.

Just as computer hardware follows Moore’s Law, which says that computing power doubles every eighteen months, so too does computer capacity and functionality. Consider, for instance, the process of legal discovery. As Carr describes it,

computers can [now] parse thousands of pages of digitized documents in seconds. Using e-discovery software with language-analysis algorithms, the machines not only spot relevant words and phrases but also discern chains of events, relationships among people, and even personal emotions and motivations. A single computer can take over the work of dozens of well-paid professionals.



Israel Museum, Jerusalem/Bridgeman Images

'Bizarre Figures'; etching by Giovanni Battista Bracelli, 1624

Or take the autonomous automobile. It can sense all the vehicles around it, respond to traffic controls and sudden movements, apply the brakes as needed, know when the tires need air, signal a turn, and never get a speeding ticket. Volvo predicts that by 2020 its vehicles will be “crash-free,” but even now there are cars that can park themselves with great precision.

The goal of automating automobile parking, and of automating driving itself, is no different than the goal of automating a factory, or pharmaceutical discovery, or surgery: it's to rationalize the process, making it more efficient, productive, and cost-effective. What this means is that automation is always going to be more convenient than what came before it—for someone. And while it's often pitched as being most convenient for the end user—the patient on the operating table, say, or the Amazon shopper, or the Google searcher, in fact the rewards of convenience flow most directly to those who own the automated system (Jeff Bezos, for example, not the Amazon Prime member).

Since replacing human labor with machine labor is not simply the collateral damage of automation but, rather, the point of it, whenever the workforce is subject to automation, technological unemployment, whether short- or long-lived, must follow. The MIT economists Eric Brynjolfsson and Andrew McAfee, who are champions of automation, state this unambiguously when they write:

Even the most beneficial developments have unpleasant consequences that must be managed.... Technological progress is going to leave behind some people, perhaps even a lot of people, as it races ahead.¹

Flip this statement around, and what Brynjolfsson and McAfee are also saying is that while technological progress is going to force many people to submit to tightly monitored control of their movements, with their productivity clearly measured, that progress is also going to benefit perhaps just a few as it races ahead. And that, it appears, is what is happening. (Of the fifteen wealthiest Americans, six own digital technology companies, the oldest of which, Microsoft, has been in existence only since 1975. Six others are members of a single family, the Waltons, whose vast retail empire, with its notoriously low wages, has meant that people are much cheaper and more expendable than warehouse robots. Still, Walmart has benefited from an automated point-of-sale system that enables its owners to know precisely what is selling where and when, which in turn allows them to avoid stocking slow-moving items and to tie up less money than the competitors in inventory.)

As Paul Krugman wrote a couple of years ago in *The New York Times*:

Smart machines may make higher GDP possible, but they will also reduce the demand for people—including smart people. So we could be looking at a society that grows ever richer, but in which all the gains in wealth accrue to whoever owns the robots.

In the United States, real wages have been stagnant for the past four decades, while corporate profits have soared. As of last year, 16 percent of men between eighteen and fifty-four and 30 percent of women in the same age group were not working, and more than a third of those who were unemployed attributed their joblessness to technology. As *The Economist* reported in early 2014:

Recent research suggests that...substituting capital for labor through automation is increasingly attractive; as a result owners of capital have captured ever more of the world's income since the 1980s, while the share going to labor has fallen.

There is a certain school of thought, championed primarily by those such as Google's Larry Page, who stand to make a lot of money from the ongoing digitization and automation of just about everything, that the elimination of jobs concurrent with a rise in productivity will lead to a leisure class freed from work. Leaving aside questions about how these lucky folks will house and feed themselves, the belief that most people would like nothing more than to be able to spend all day in their pajamas watching TV—which turns out to be what many “nonemployed” men do—sorely misconstrues the value of work, even work that might appear to an outsider to be less than fulfilling. Stated simply: work confers identity. When Dublin City University professor Michael Doherty surveyed Irish workers, including those who stocked grocery shelves and drove city buses, to find out if work continues to be “a significant locus of personal identity,” even at a time when employment itself is less secure, he concluded that “the findings of this research can be summed up in the succinct phrase: ‘work matters.’”²

How much it matters may not be quantifiable, but in an essay in *The New York Times*, Dean Baker, the codirector of the Center for Economic and Policy Research, noted that there was

a 50 to 100 percent increase in death rates for older male workers in the years immediately following a job loss, if they previously had been consistently employed.

One reason was suggested in a study by Mihaly Csikszentmihalyi, the author of *Flow: The Psychology of Optimal Experience* (1990), who found, Carr reports, that “people were

happier, felt more fulfilled by what they were doing, while they were at work than during their leisure hours.”

Even where automation does not eliminate jobs, it often changes the nature of work. Carr makes a convincing case for the ways in which automation dulls the brain, removing the need to pay attention or master complicated routines or think creatively and react quickly. Those airline pilots who now are at the controls for less than three minutes find themselves spending most of their flight time staring at computer screens while automated systems do the actual flying. As a consequence, their overreliance on automation, and on a tendency to trust computer data even in the face of contradictory physical evidence, can be dangerous. Carr cites a study by Matthew Ebbatson, a human factors researcher, that

found a direct correlation between a pilot’s aptitude at the controls and the amount of time the pilot had spent flying without the aid of automation.... The analysis indicated that “manual flying skills decay quite rapidly towards the fringes of ‘tolerable’ performance without relatively frequent practice.”

Similarly, an FAA report on cockpit automation released in 2013 found that over half of all airplane accidents were the result of the mental autopilot brought on by actual autopilot.

If aviation is a less convincing case, since the overall result of automation has been to make flying safer, consider a more mundane and ubiquitous activity, Internet searches using Google. According to Carr, relying on the Internet for facts and figures is making us mindless sloths. He points to a study in *Science* that demonstrates that the wealth of information readily available on the Internet disinclines users from remembering what they’ve found out. He also cites an interview with Amit Singhal, Google’s lead search engineer, who states that “the more accurate the machine gets [at predicting search terms], the lazier the questions become.”

A corollary to all this intellectual laziness and dullness is what Carr calls “deskilling”—the loss of abilities and proficiencies as more and more authority is handed over to machines. Doctors who cede authority to machines to read X-rays and make diagnoses, architects who rely increasingly on computer-assisted design (CAD) programs, marketers who place ads based on algorithms, traders who no longer trade—all suffer a diminution of the expertise that comes with experience, or they never gain that experience in the first place. As Carr sees it:

As more skills are built into the machine, it assumes more control over the work, and the worker’s opportunity to engage in and develop deeper talents, such as those involved in interpretation and judgment, dwindles. When automation reaches its highest level, when it takes command of the job, the worker, skillwise, has nowhere to go but down.

Conversely, machines have nowhere to go but up. In Carr’s estimation, “as we grow more reliant on applications and algorithms, we become less capable of acting without their aid.... That makes the software more indispensable still. Automation breeds automation.”

But since automation also produces quicker drug development, safer highways, more accurate medical diagnoses, cheaper material goods, and greater energy efficiency, to name just a few of its obvious benefits, there have been few cautionary voices like Nicholas Carr’s urging us to take stock, especially, of the effects of automation on our very

humanness—what makes us who we are as individuals—and on our humanity—what makes us who we are in aggregate. Yet shortly after *The Glass Cage* was published, a group of more than one hundred Silicon Valley luminaries, led by Tesla’s Elon Musk, and scientists, including the theoretical physicist Stephen Hawking, issued a call to conscience for those working on automation’s holy grail, artificial intelligence, lest they, in Musk’s words, “summon the demon.” (In Hawking’s estimation, AI could spell the end of the human race as machines evolve faster than people and overtake us.) Their letter is worth quoting at length, because it demonstrates both the hubris of those who are programming our future and the possibility that without some kind of oversight, the golem, not God, might emerge from their machines:

[Artificial intelligence] has yielded remarkable successes in various component tasks such as speech recognition, image classification, autonomous vehicles, machine translation, legged locomotion, and question-answering systems.

As capabilities in these areas and others cross the threshold from laboratory research to economically valuable technologies, a virtuous cycle takes hold whereby even small improvements in performance are worth large sums of money, prompting greater investments in research....

The potential benefits are huge, since everything that civilization has to offer is a product of human intelligence; we cannot predict what we might achieve when this intelligence is magnified by the tools AI may provide, but the eradication of disease and poverty are not unfathomable. Because of the great potential of AI, it is important to research how to reap its benefits while avoiding potential pitfalls.

The progress in AI research makes it timely to focus research not only on making AI more capable, but also on maximizing the societal benefit.... [Until now the field of AI] has focused largely on techniques that are neutral with respect to purpose. We recommend expanded research aimed at ensuring that increasingly capable AI systems are robust and beneficial: our AI systems must do what we want them to do.

Just who is this “we” who must ensure that robots, algorithms, and intelligent machines act in the public interest? It is not, as Nicholas Carr suggests it should be, the public. Rather, according to the authors of the research plan that accompanies the letter signed by Musk, Hawking, and the others, making artificial intelligence “robust and beneficial,” like making artificial intelligence itself, is an engineering problem, to be solved by engineers. To be fair, no one but those designing these systems is in a position to build in measures of control and security, but what those measures are, and what they aim to accomplish, is something else again. Indeed, their research plan, for example, looks to “maximize the economic benefits of artificial intelligence while mitigating adverse effects, which could include increased inequality and unemployment.”

The priorities are clear: money first, people second. Or consider this semantic dodge: “If, as some organizations have suggested, autonomous weapons should be banned, is it possible to develop a precise definition of autonomy for this purpose...?” Moreover, the authors acknowledge that “aligning the values of powerful AI systems with our own values and preferences [may be] difficult,” though this might be solved by building “systems that can learn or acquire values at run-time.” However well-meaning, they fail to say what values, or whose, or to recognize that most values are not universal but, rather, culturally and socially constructed, subjective, and inherently biased.

We live in a technophilic age. We love our digital devices and all that they can do for us. We celebrate our Internet billionaires: they show us the way and deliver us to our destiny. We have President Obama, who established the National Robotics Initiative to develop the “next generation of robotics, to advance the capability and usability of such systems and artifacts, and to encourage existing and new communities to focus on innovative application areas.” Even so, it is naive to believe that government is competent, let alone in a position, to control the development and deployment of robots, self-generating algorithms, and artificial intelligence. Government has too many constituent parts that have their own, sometimes competing, visions of the technological future. Business, of course, is self-interested and resists regulation. We, the people, are on our own here—though if the AI developers have their way, not for long.

* Carr discusses integrated development environments (IDEs) which programmers use to check their code, and quotes Vivek Haldar, a veteran Google developer: “‘The behavior all these tools encourage is not ‘think deeply about your code and write it carefully,’ but ‘just write a crappy first draft of your code, and then the tools will tell you not just what’s wrong with it, but also how to make it better.’” [↪](#)

1 *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (Norton, 2014), pp. 10–11. [↪](#)

2 Michael Doherty, “When the Working Day Is Through: The End of Work As Identity?” *Work, Employment and Society*, Vol. 23, No. 1 (March 2009). [↪](#)

RELATED



Bot or Not?
[James Gleick](#)



The Creepy New Wave of the Internet
[Sue Halpern](#)



Why Work Is More and More Debased [🔒](#)
[Robert Kuttner](#)