

CHAPTER ONE

COMPUTERS ARE IMPROVING FASTER THAN YOU ARE

As Technology Becomes More Awesomely Able,
What Will Be the High-Value Human Skills of Tomorrow?

I am standing on a stage, behind a waist-high podium with my first name on it. To my right is a woman named Vicki; she's behind an identical podium with her name on it. Between us is a third podium with no one behind it, just the name "Watson" on the front. We are about to play *Jeopardy!*

This is the National Retail Federation's mammoth annual conference at New York City's Javits Center, and in addition to doing some onstage moderating, I have insanely agreed to compete against IBM's Watson, the cognitive computing system, whose power the company wants to demonstrate to the 1,200 global retail leaders sitting in front of me. Watson's celebrated defeat of *Jeopardy!*'s two greatest champions is almost a year old, so I'm not expecting this to go well. But I'm not prepared for what hits me.

We get to a category called "Before and After at the Movies." *Jeopardy!* aficionados have seen this category many times over the years, but I have never heard of it. First clue, for \$200: "Han

Solo meets up with Lando Calrissian while time traveling with Marty McFly.”

Umm . . . what?

Watson has already buzzed in. “What is *The Empire Strikes Back to the Future?*” it responds correctly.

It picks the same category for \$400: “James Bond fights the Soviets while trying to romance Ali MacGraw before she dies.” I’m still struggling with the concept, but Watson has already buzzed in. “What is *From Russia with Love Story?*” Right again.

By the time I figure this out, Watson is on the category’s last clue: “John Belushi & the boys set up their fraternity in the museum where crazy Vincent Price turns people into figurines.” The correct response, as Watson instantly knows, is “What is *Animal House of Wax?*” Watson has run the category.

My humiliation is not totally unrelieved. I do get some questions right in other categories, and Watson gets some wrong. But at the end of our one round I have been shellacked. I actually don’t remember the score, which must be how the psyche protects itself. I just know for sure that I have witnessed something profound.

Realize that Watson is not connected to the Internet. It’s a free-standing machine just like me, relying only on what it knows. It has been loaded with the entire contents of Wikipedia, for example, and much, much more. No one types the clues into Watson; it has to hear and understand the emcee’s spoken words, just as I do. In addition, Watson is intentionally slowed down by a built-in delay when buzzing in to answer a clue. We humans must use our prehistoric muscle systems to push a button that closes a circuit and sounds the buzzer. Watson could do it at light speed with an electronic signal, so the developers interposed a delay to level the playing field. Otherwise I’d never have a prayer of winning, even if we both knew the correct response. But, of course, even with the delay, I lost.

So let’s confront reality: Watson is smarter than I am. In fact,

I’m surrounded by technology that’s better than I am at sophisticated tasks. Google’s autonomous car is a better driver than I am. The company has a whole fleet of vehicles that have driven hundreds of thousands of miles with only one accident while in autonomous mode, when one of the cars was rear-ended by a human driver at a stoplight. Computers are better than humans at screening documents for relevance in the discovery phase of litigation, an activity for which young lawyers used to bill at an impressive hourly rate. Computers are better at detecting some kinds of human emotion, despite our million years of evolution that was supposed to make us razor sharp at that skill.

One more thing. I competed against Watson in early 2012. Back then it was the size of a bedroom. As I write, it has shrunk to the size of three stacked pizza boxes, yet it’s also 2,400 percent faster.

More broadly, information technology is doubling in power roughly every two years. I am not—and I’ll guess that you’re not either.

A NIGHTMARE FUTURE?

The mind-bending progress of information technology makes it easier every day for us to imagine a nightmare future. Computers become so capable that they’re simply better at doing thousands of tasks that people now get paid to do. Sure, we’ll still need people to make high-level decisions and to develop even smarter computers, but we won’t need enough such workers to keep the broad mass of working-age people employed, or for their living standard to rise. And so, in the imaginary nightmare future, millions of people will lose out, unable finally to best the machine, struggling hopelessly to live the lives they thought they had earned.

In fact, as we shall see, substantial evidence suggests that

technology advances really are playing a role in increasingly stubborn unemployment, slow wage growth, and the trend of college graduates taking jobs that don't require a bachelor's degree. If technology is actually a significant cause of those trends, then the miserable outlook becomes hard to dismiss.

But that nightmare future is not inevitable. Some people have suffered as technology has taken away their jobs, and more will do so. But we don't need to suffer. The essential reality to grasp, larger than we may realize, is that the very nature of work is changing, and the skills that the economy values are changing. We've been through these historic shifts a few times before, most famously in the Industrial Revolution. Each time, those who didn't recognize the shift, or refused to accept it, got left behind. But those who embraced it gained at least the chance to lead far better lives. That's happening this time as well.

While we've seen the general phenomenon before, the way that work changes is different every time, and this time the changes are greater than ever. The skills that will prove most valuable are no longer the technical, classroom-taught, left-brain skills that economic advances have demanded from workers over the past 300 years. Those skills will remain vitally important, but important isn't the same as valuable; they are becoming commoditized and thus a diminishing source of competitive advantage. The new high-value skills are instead part of our deepest nature, the abilities that literally define us as humans: sensing the thoughts and feelings of others, working productively in groups, building relationships, solving problems together, expressing ourselves with greater power than logic can ever achieve. These are fundamentally different types of skills than those the economy has valued most highly in the past. And unlike some previous revolutions in what the economy values, this one holds the promise of making our work lives not only rewarding financially, but also richer and more satisfying emotionally.

Step one in reaching that future is to think about it in a new way. We shouldn't focus on beating computers at what they do. We'll lose that contest. Nor should we even follow the inviting path of trying to divine what computers inherently cannot do—because they can do more every day.

The relentless advance of computer capability is of course merely Moore's Law at work, as it has been for decades. Still, it's hard for us to appreciate all the implications of this simple trend. That's because most things in our world slow down as they get bigger and older; for evidence, just look in the mirror. It's the same with other living things, singly or in groups. From protozoa to whales, everything eventually stops growing. So do organizations. A small start-up company can easily grow 100 percent a year, but a major Fortune 500 firm may struggle to grow 5 percent.

Technology isn't constrained that way. It just keeps getting more powerful. Sony's first transistor radio was advertised as pocket-sized, but it was actually too big, so the company had salesmen's shirts specially made with extra-large pockets; that radio had five transistors. Intel's latest chip, the size of your thumbnail, has five billion transistors, and its replacement will have ten billion. Today's infotech systems, having become as awesomely powerful as they are, will be 100 percent more awesomely powerful in two years. Moore's law must end eventually, but new technologies in development could be just as effective, and better algorithms are already multiplying computing power in some cases even more than hardware improvements are doing. To imagine that technology won't keep advancing at a blistering pace seems unwise.

Consider what is being doubled. It isn't just year-before-last's achievement in computing power. What gets doubled every two years is everything that has been achieved in the history of computing power up to that point. Back when that progression meant going from five transistors in a device to ten, it didn't much change the

world. Now that it means going from five billion transistors on a tiny chip to ten billion to twenty billion to forty billion—that's three doublings, just six years—it means literally more than we can imagine.

That's because it's so unlike everything else in our world in ways even beyond physical growth rates. For us humans, learning, like growing, gets harder with time. When humans learn to do something, we make slow progress at first—learning how to hold the golf club or turn the steering wheel smoothly—then rapid progress as we get the hang of it, and then our advancement slows down. Pretty soon, most of us are as good as we're going to get. We can certainly keep improving through devoted practice, but each advance is typically a bit smaller than the one before.

Information technology is just the opposite. When a doubling of computing power for a given price meant going from five transistors to ten, it made a device smarter by only five transistors. Now, after many doublings, the current doubling will make a device smarter by five billion transistors, and the next one will make a device smarter by ten billion.

While people get more skilled by ever smaller increments, computers get more capable by ever larger ones.

The issue is clear and momentous. As technology becomes more capable, advancing inexorably by ever longer two-year strides and acquiring abilities that are increasingly complex and difficult, what will be the high-value human skills of tomorrow—the jobs that will pay well for us and our kids, the competencies that will distinguish winning companies, the traits of dominant nations? To put it starkly: What will people do better than computers?

CHAPTER TWO

GAUGING THE CHALLENGE

A Growing Army of Experts Wonder If Just Maybe
the Luddites Aren't Wrong Anymore.

In the movie *Desk Set*, a 1957 romantic comedy starring Katharine Hepburn and Spencer Tracy, Hepburn plays the head of the research department at a major TV network. Today a TV network's research department focuses entirely on audience research, but back then it was a general information resource for anyone at the company, and yes, networks and other companies really had such departments. Equipped with two floors of reference works and other books, its staff stood ready to supply any information that any employee might ask for—the opening lines of *Hiawatha*, the weight of the earth, the names of Santa's reindeer (all of which were queries for Hepburn's department in the movie). That is, employees could pick up the phone, call Katharine Hepburn's character, and ask in their own words for any information, and she and her staff would search a vast trove of data and return an answer far faster than the caller could ever have found it.

Hepburn's character is named Miss Watson.

All is well until one day the network boss decides to install a

computer—an “electronic brain,” they call it—named EMERAC (a clear reference to ENIAC and UNIVAC, the wonder machines of the era). It was invented by the Spencer Tracy character. Shortly before Miss Watson hears the news that EMERAC is coming to her department, she sees it demonstrated elsewhere, translating Russian into Chinese, among other feats. Her assessment, as expressed to her coworkers: “Frightening. Gave me the feeling that maybe, just maybe, people were a little bit outmoded.”

The Tracy character, Richard Sumner, shows up to install the machine, and Miss Watson and her staff assume they’ll be fired once it’s up and running. In a memorable scene, he demonstrates the machine to a group of network executives and explains its advantages:

Sumner: “The purpose of this machine, of course, is to free the worker—”

Miss Watson: “You can say that again.”

Sumner: “—to free the worker from the routine and repetitive tasks and liberate his time for more important work.”

Miss Watson and the rest of the research staff are indeed fired, but before they can clean out their desks, EMERAC botches some requests it can’t handle—a call for information on Corfu, for example, returns reams of useless data on the word “curfew,” while a staffer scurries into the stacks and gets the needed answers the old-fashioned way. And then it turns out that the researchers actually should not have received termination notices after all. An EMERAC computer in the payroll department had gone haywire and fired everyone in the company. The error is corrected, the research staffers keep their jobs and learn to work with EMERAC, Miss Watson wisely decides to marry the Spencer Tracy character and not the Gig Young character (part of the mandatory romantic subplot), and once again all is well.

Desk Set is extraordinarily prescient about some future capabilities and uses of computers, and also faithful to the fearful popular sentiment about them. Of course, Miss Watson is exactly the human predecessor of today’s Watson cognitive computing system. (Are the names a coincidence? The film’s opening credits include this intriguing one: “The filmmakers gratefully acknowledge the cooperation and assistance of the International Business Machines Corporation.” IBM’s founder was Thomas J. Watson, namesake of today’s Watson computing system, and his son was CEO at the time of the film.) EMERAC as explained by Sumner in the film is remarkably similar to today’s Watson: All the information in all those books in the research library—encyclopedias, atlases, Shakespeare’s plays—was fed into the machine, which could then respond instantly to natural-language requests (typed, not spoken) for information. Even in 1957 the idea was clear; the technology just wasn’t ready.

The research staffers’ fears about being replaced by a computer were also a sign of things to come. “I hear thousands of people are losing their jobs to these electronic brains,” one of them says. She heard right, and the thousands would become millions. At the same time, the corporate response intended to calm those fears has remained just what Sumner said in the movie—that computers would “free the worker from the routine and repetitive tasks” so he or she could do “more important work.” To this day it’s striking how everyone working on advanced information technology seems to feel defensive about the implicit threat of eliminating jobs and takes pains to say that they’re not trying to replace people. “We’re not intending to replace humans,” said Kirstin Petersen of Harvard’s Wyss Institute for Biologically Inspired Engineering, in explaining the institute’s development of “swarm robotics,” in which large numbers of small, simple robots do construction jobs. “We’re intending to work in situations where humans can’t work or it’s impractical for them to work.” IBM

has always said that Watson is intended to supplement human decision making, not replace it—"to make people more intelligent about what they do."

Most important, and perhaps surprising, is that even the film's happily-ever-after ending was realistic in the large sense, at least with regard to employment, if not romance. Viewed on the scale of the entire economy, technology's advance indeed has not cost jobs, despite the widespread fears. Quite the opposite. And those fears are much more deeply rooted than most of us realize.

THE NEW SKEPTICS

The conventional view is that fear of technology arose when technology started upending the economic order in the eighteenth century, at the start of the Industrial Revolution in Britain. But the fears were already well entrenched, and innovators were already sounding remarkably modern in arguing that technology was a boon, not a bane, for workers. In the late sixteenth century, an English clergyman named William Lee invented a machine for knitting stockings—a wonderful advance, he believed, because it would liberate hand knitters from their drudgery. When he demonstrated it to Queen Elizabeth I in 1590 or so and asked for a patent, she reportedly replied, "Thou aimest high, Master Lee. Consider thou what the invention could do to my poor subjects. It would assuredly bring them to ruin by depriving them of employment, thus making them beggars." After the royal slap down, the queen denied his patent, the hosiers' guild campaigned against him, and he was forced to move to France, where he died in poverty.

Some 150 years later, in the early dawn of the Industrial Revolution, an Englishman named John Kay revolutionized weaving by inventing the flying shuttle, which doubled productivity—surely a boon for weavers, who could now make twice as much cloth. Yet

weavers campaigned against him, manufacturers conspired to violate his patents, and he was forced to move to France, where he died in poverty just like William Lee. Dying destitute in France seemed to be an occupational hazard for innovators.

By the time the Industrial Revolution got going, the pattern was well established. People hated technology that improved productivity. Luddites, smashing power looms in the early nineteenth century, were only the most famous exemplars.

These protesters were right in the short run, but in the long run they were resoundingly wrong. New technology does destroy jobs, but it also creates new ones—jobs for people who operate the stocking frames and power looms, for example. More important, better technology creates better jobs. Workers using improved technology are more productive, so they earn more—and spend more, creating more new jobs across the economy. At the same time, the products those tech-enabled workers make cost less than before; machine-made cloth costs a fraction of what the handmade version costs. The result is that technology, over time and across economies, has raised living standards spectacularly. For centuries, the fears of Luddites past and present have been not merely unfounded but the exact opposite of reality. Advancing technology has improved the material well-being of humanity more than any other development in history, by far.

Now something has changed. The way technology benefits workers is one of the firmest orthodoxies in all of economics, but recently, for the first time, many mainstream economists and technologists have begun to question whether it will continue.

The proximate cause of their new skepticism is the sorry job-generating performance of the developed economies in the wake of the 2008–2009 financial crisis and recession. For decades, the U.S. economy regularly returned to prerecession employment levels about eighteen months after a recession started. Then, starting

with the 1990–1991 recession, the lag started lengthening. After the 2008–2009 recession the recovery of employment took seventy-seven months—over six years. How come? And why did wages begin stagnating for large swaths of the U.S. workforce long before the recession began? Why is the same trend happening in other developed countries? As economists look for answers, they see factors that go far beyond the causes of the recession.

“THE DEFINING ECONOMIC FEATURE OF OUR ERA”

Lawrence H. Summers—former U.S. treasury secretary, former president of Harvard University, a star economist—is one of the new skeptics. In a significant lecture to an audience of fellow economists, he summarized in his brisk way the orthodox view of the debate over technology: “There were the stupid Luddite people, who mostly were outside of economics departments, and there were the smart progressive people. . . . The stupid people thought that automation was going to make all the jobs go away and there wasn’t going to be any work to do. And the smart people understood that when more was produced, there would be more income and therefore there would be more demand. It wasn’t possible that all the jobs would go away, so automation was a blessing.”

Evidence overwhelmingly supported that view for decades. All you had to do was imagine the world of 1800 and compare it with the world around you. But then, quite recently, the world changed: “Until a few years ago, I didn’t think this was a very complicated subject,” Summers said. “The Luddites were wrong and the believers in technology and technological progress were right. I’m not so completely certain now.”

Summers is far from the only expert who became doubtful. The Pew Research Center Internet Project in 2014 canvassed 1,896 experts it had identified as insightful on technology issues, and it

asked them this question: Will technology displace more jobs than it creates by 2025? Half said yes, and half said no. That was an astounding result. As Summers explained, the evidence in favor of “no” was perfectly clear, or it had been. It’s hard to imagine that, ten years before, as many as 10 percent of such a highly informed group would have said yes. (We don’t know for sure because apparently no one thought the question even worth asking.) Now half said so. The orthodoxy was suddenly no longer orthodox.

What Summers and other economists believe has changed is, in concept, simple. The two factors of production are capital and labor, and in economists’ terms they have always been regarded as complements, not just substitutes. Capital makes workers more productive. Even if it displaces some workers (substitutes for them), it also creates new, more productive jobs using that new capital so that, as Summers said, “if there’s more capital, the wage has to rise” (it complements workers). But now he and others began seeing a new possibility: Capital can substitute for labor, period. Summers explained, “That is, you can take some of the stock of machines and, by designing them appropriately, you can have them do exactly what labor did before.”

The key word is “exactly.” A Google self-driving car doesn’t complement anybody’s work because nobody operates it at all. The company produced a version that doesn’t have a steering wheel, brake pedal, or accelerator, and it’s designed to transport even blind or other disabled people. So it doesn’t make drivers, even a shrunken population of them, more productive. It does exactly what they do and thus just replaces them.

In a world like that, economic logic dictates that wage rates must fall, and the share of total income going to capital rather than labor must rise, which is indeed what has been happening. An important reason, Summers says, is “the nature of the technical changes that we have seen: Increasingly they take the form of capital that effectively substitutes for labor.”

The outlook is obviously for much more capital-labor substitution as computing power gallops unflaggingly forward. That is not a happy future for many people. In fact, as Summers reasons, "It may well be that, given the possibilities for substitution, some categories of labor will not be able to earn a subsistence income."

Economists aren't the only experts who see such a trend. "Unlike previous disruptions, such as when farming machinery displaced farm workers but created factory jobs making the machines, robotics and AI [artificial intelligence] are different," Mark Nall, a NASA program manager with much real-world technology experience, told the Pew canvassers. "Due to their versatility and growing capabilities, not just a few economic sectors will be affected, but whole swaths will be. . . . The social consequence is that good-paying jobs will be increasingly scarce." Stowe Boyd, lead researcher at Gigaom Research, a technology research firm, was even more pessimistic: "An increasing proportion of the world's population will be outside the world of work—either living on the dole, or benefiting from the dramatically decreased costs of goods to eke out a subsistence lifestyle." Michael Roberts, a much respected Internet pioneer, predicted confidently that "electronic human avatars with substantial work capability are years, not decades away. . . . There is great pain down the road for everyone as new realities are addressed. The only question is how soon."

Microsoft founder Bill Gates has observed the trend also and believes it's greatly underappreciated: "Software substitution, whether it's for drivers or waiters or nurses—it's progressing," he told a Washington, D.C., audience in 2014. "Technology over time will reduce demand for jobs. . . . Twenty years from now, labor demand for lots of skill sets will be substantially lower. I don't think people have that in their mental model."

But isn't all this garment rending and teeth gnashing just the usual worry over the endless cycle of creative destruction, as new

industries displace old ones? You can't earn a subsistence income with the skills of making slide rules, and that's not a problem because you can earn a better income doing something else. But the analogy isn't valid. You can't earn a living by making slide rules because nobody wants them anymore. This new argument, by contrast, holds that the economy can increasingly provide exactly the goods and services that people most want today and tomorrow, and can do it using more machines and ever fewer people.

Thus Summers's conclusion, which is significant coming from an economist of his stature: "This set of developments is going to be the defining economic feature of our era."

THE FOURTH GREAT TURNING POINT FOR WORKERS

The immediate question for most of us is obvious: Who, specifically, gets hurt, and who doesn't?

To find the answer, it helps to think of these developments as the latest chapter in a story. Technology has been changing the nature of work and the value of particular skills for well over 200 years, and the story so far comprises just three major turning points.

At first, the rise of industrial technology devalued the skills of artisans, who handcrafted their products from beginning to end: A gun maker carved the stock, cast the barrel, engraved the lock, filed the trigger, and painstakingly fitted the pieces together. But in Eli Whitney's Connecticut gun factory, separate workers did each of those jobs, or just portions of them, using water-powered machinery, and components of each type were identical. Skilled artisans were out of luck, but less skilled workers were in demand. They could easily learn to use the new machines—the workers and machines were complements—and so the workers could earn far more than before.

The second turning point arrived in the early twentieth century,

when a new trend emerged. Widely available electricity enabled the building of far more sophisticated factories, requiring better educated, more highly skilled workers to operate the more complicated machines; companies also grew much larger, requiring a larger corps of educated managers. Now the unskilled were out of luck, and educated workers were in demand—but that was okay, because the unskilled could get educated. The trend intensified through most of the twentieth century. Advancing technology continually required better educated workers, and Americans responded by educating themselves with unprecedented ambition. The high school graduation rate rocketed from 4 percent in 1890 to 77 percent in 1970, a national intellectual upgrade such as the world had never seen. As long as workers could keep up with the increasing demands of technology, the two remained complements. The result was an economic miracle of fast-rising living standards.

But then the third major turning point arrived, starting in the 1980s. Information technology had developed to a point where it could take over many medium-skilled jobs—bookkeeping, back-office jobs, repetitive factory work. The number of jobs in those categories diminished, and wages stagnated for the shrinking group of workers who still did them. Yet the trend was limited. At both ends of the skill spectrum, people in high-skill jobs and low-skill service jobs did much better. The number of jobs in those categories increased, and pay went up. Economists called it the polarization of the labor market, and they observed it in the United States and many other developed countries. At the top end of the market, infotech still wasn't good enough to take over the problem-solving, judging, and coordinating tasks of high-skill workers like managers, lawyers, consultants, and financiers; in fact, it made those workers more productive by giving them more information at lower cost. At the bottom end, infotech didn't threaten low-skill service workers because computers were terrible at skills of physical dexterity; a computer could

defeat a grand master chess champion but couldn't pick up a pencil from a tabletop. Home health aides, gardeners, cooks, and others could breathe easy.

That was the story into the 2000s. In the nonstop valuing and devaluing of skills through economic history, infotech was crushing medium-skill workers, but workers at the two ends of the skill spectrum were safe or prospering. Now we are at a fourth turning point. Infotech is advancing steadily into both ends of the spectrum, threatening workers who thought they didn't have to worry.

MAYBE EVEN LAWYERS CAN'T OUTSMART COMPUTERS

At the top end, what's happening to lawyers is a model for any occupation involving analysis, subtle interpretation, strategizing, and persuasion. The computer incursion into the legal discovery process is well known. In cases around the world computers are reading millions of documents and sorting them for relevance without ever getting tired or distracted. The cost savings are extraordinary. One e-discovery vendor, Symantec's Clearwell, claimed it could cut costs up to 98 percent. That may seem outlandish, but it's in line with the claims of an executive at another vendor, Autonomy, who told the *New York Times* that e-discovery would enable one lawyer to do the work of 500 or more. In addition, software does the job much, much better than people. It can detect patterns in thousands or millions of documents that no human could spot—unusual editing of a document, for example, or spikes in communication between certain people, or even changes in e-mail style that may signal hidden motives.

But that's just the beginning. Computers then started moving up the ladder of value, becoming highly skilled at searching the legal literature for appropriate precedents in a given case, and doing it far more widely and thoroughly than people can do. Humans still have to identify the legal issues involved, but, as Northwestern University

law professor John O. McGinnis has written, “search engines will eventually do this by themselves, and then go on to suggest the case law that is likely to prove relevant to the matter.”

Advancing even higher into the realm of lawyerly skill, computers can already predict Supreme Court decisions better than legal experts can. As such analytical power expands in scope, computers will move nearer the heart of what lawyers do by advising better than lawyers can on whether to sue or settle or go to trial before any court and in any type of case. Companies such as Lex Machina and Huron Legal already offer such analytic services, which are improving by the day. These firms’ computers have read all the documents in hundreds of thousands of cases and can tell you, for example, which companies are more likely to settle than to litigate a patent case, or how particular judges tend to rule in particular types of cases, or which lawyers have the best records in front of specified judges. As more potential litigants, both plaintiffs and defendants, can see better analysis of vastly more data, odds are strong they’ll be able to resolve disputes far more efficiently. One possible result: fewer lawsuits.

None of this means that lawyers will disappear, but it suggests that the world will need fewer of them. It’s already happening. “The rise of machine intelligence is probably partly to blame for the current crisis of law schools”—shrinking enrollments, falling tuitions—“and will certainly worsen that crisis,” McGinnis has observed.

With infotech thoroughly disrupting even a field so advanced that it requires three years of post-graduate education and can pay extremely well, other high-skill workers—analysts, managers—can’t help but wonder about their own futures. What’s happening in law is the application of Watson-like technology to a specific industry, but it can be applied far more widely. The breakthrough of this technology is that it understands natural language, so when you ask it a question, it doesn’t just search for keywords from the question you asked. It tries to figure out the context of your question and thus understand

what you really mean. So, for example, if your question includes the phrase “two plus two,” that might mean “four,” or, if you’re in the car business, it might mean “a car with two front seats and two back seats,” or if you’re a psychologist, it might mean “a family with two parents and two children.” Cognitive computing systems derive the context and then come up with possible answers to your question and estimate which one is most likely correct. A system’s answers aren’t especially good when it first delves into a field, but with experience it keeps getting better. That’s why the Internet entrepreneur Terry Jones, who founded Travelocity, has said that “Watson is the only computer that’s worth more used than new.”

Watson-like technology works best when it has a really big body of written material to read and work with. For *Jeopardy!* Watson downloaded not only the entire contents of Wikipedia but also thousands of previous *Jeopardy!* clues and responses. Law is obviously an excellent field for this technology. Medicine is another. Memorial Sloan Kettering Cancer Center in New York City uses Watson to extract answers from the vast oncology literature, a task which no physician could ever keep up with. Financial advice looks like a fat target for this technology because it involves a vast and growing corpus of research plus huge volumes of data that change every day. Several financial institutions are therefore using Watson, initially as a tool for their financial advisers. But look just a bit down the road: Corporate Insight, a research firm that focuses on the financial services industry, asks, “Once consumers have a personal Watson in their pocket . . . why would an experienced investor need a financial adviser?”

WRITERS WHO NEVER GET BLOCKED, TIRED, OR DRUNK

Combine an understanding of natural language with high-torque analytic power and you get a nonfiction writer, or at least a species of

one. A company called Narrative Science makes software that writes articles that would not strike most people as computer-written. It focused first on events embodying lots of data: ball games and corporate earnings announcements. The software became increasingly sophisticated at going beyond the facts and figures—for example, figuring out the most important play in a game or identifying the best angle for the article: a come-from-behind win, say, or a hero player. Then the developers taught the software different writing styles, which customers could choose from a menu. Next, it learned to understand more than just numerical data, reading relevant material to create context for the article. A number of media companies, including Yahoo and Forbes, publish articles from Narrative Science, though some of the company's customers don't want to be identified and don't tell readers which articles are computer-written. In mid-2014, the Associated Press assigned computers to write all its articles about corporate earnings announcements.

Then Narrative Science realized that maybe the real money wasn't in producing journalism at all (they could have asked any journalist about that) but in generating the writing that companies use internally, the countless reports and analyses that influence business decisions. So it arranged its technology to gather broad classes of data, including unstructured data like social media posts, on any given topic or problem, and to analyze it deeply, looking for trends, correlations, unusual events, and more. The software uses that data to "make judgments and draw conclusions," the company says; it can also make recommendations. The software writes it all up at a reading level and in a tone that the customer chooses, also supplying helpful charts and graphs.

This is starting to sound less like writing and more like management.

But are the writing and the analysis any good? That, at least, is for humans to decide. Except that increasingly it need not be. Schools

from the elementary level through college are using software to judge writing and analysis in the form of student essays. The software isn't perfect—it doesn't yet evaluate such subtleties as voice and tone—but human graders aren't perfect either. Jeff Pence, a middle-school teacher in Canton, Georgia, who used the software to help grade papers from his 140 students, acknowledged that it doesn't grade with perfect accuracy, but, he told *Education Week*, "When I reach that 67th essay, I'm not real accurate." Similar software is being used at much higher levels. EdX, the enterprise started by Harvard and MIT to offer online courses, has begun using it to grade student papers. The Hewlett Foundation offered two \$100,000 prizes for developing such software, and edX hired one of the winners to work on its version, which is available to developers everywhere as open-source code so it can be improved.

Of course, this evaluation software must itself be evaluated by humans, measuring it against the performance of humans. So researchers had a group of human teachers grade a large set of essays. Then they gave those same essays to a separate group of human graders and to the software. They compared the grades assigned by Human Group Two to those assigned by Human Group One, and they also compared the grades assigned by the software to those assigned by Human Group One. All three sets of grades were different, but the software's grades were no more different from Human Group One's grades than Human Group Two's grades were different from Human Group One's. So while software doesn't assign the same grades as people, neither do people assign the same grades as other people. And if you look at a large group of grades assigned to the same work by people and by software, you can't tell which is which.

Two points to draw from this:

One, the software is getting rapidly better. The people are not.

Two, education as currently conceived is becoming really weird. After all, the report-writing software developed by Narrative

Science and other companies is easily adapted to other markets, such as students' papers. So we now have essay-grading software and essay-writing software, both of which are improving. What happens next is obvious. The writing software gets optimized to please the grading software. Every essay gets an A, and neither the student nor the teacher has anything to do with it. But not much education is necessarily going on, which poses problems for both student and teacher.

A ROBOT'S TOUCH

The rapid progress of infotech in taking over tasks at the high-skill end of the job spectrum—lawyers, doctors, managers, professors—is startling, but it isn't especially surprising. If we thought such jobs were by their nature immune to computer competition, we shouldn't have, because these jobs are highly cognitive. Much of the work is brain work, and that's just what computers do best; they needed only time to accumulate the required computing power. The greater surprise shows up at the opposite end of the job spectrum, in the low-skill, low-pay world where the work is less cognitive and more physical. This is the kind of work that computers for decades could hardly do at all. An example illustrates the gap in abilities: In 1997 a computer could beat the world's greatest chess player yet could not physically move the pieces on the board. But again the technology needed only time, a few more doublings of power. The skills of physical work are also not immune to the advance of infotech.

Google's autonomous cars are an obvious and significant example—significant because the number one job among American men is truck driver. Many more examples are appearing. You can train a Baxter robot (from Rethink Robotics) to do all kinds of things—pack or unpack boxes, take items to or from a conveyor belt, fold a T-shirt, carry things around, count them, inspect them—just

by moving its arms and hands ("end-effectors") in the desired way. Many previous industrial robots had to be surrounded by safety cages because they could do just one thing in one way, over and over, and that's all they knew; if you got between a welding robot and the piece it was welding, you were in deep trouble. But Baxter doesn't hurt anyone as it hums about the shop floor; it adapts its movements to its environment by sensing everything around it, including people.

Many similar kinds of robots operate in different environments—for example, buzzing through hospital hallways delivering medicines, hauling laundry, or picking up infectious waste. Security robots can hang out around public buildings, watching, listening, reading license plates, and sending information to law enforcement as the robot deems appropriate. Robots went into the wreckage of Japan's ruined Fukushima Daiichi nuclear power plant long before people did.

The advantage robots hold in doing dangerous work is a big reason the U.S. military is a major user of them and a major funder of research into them. By 2008 about 12,000 combat robots were working in Iraq. Some, barely larger than a shoebox, run on miniature tank treads and can carry a camera and other sensors; they gather intelligence and do surveillance and reconnaissance. Larger ones dispose of bombs or carry heavy loads into and out of dangerous places. A few robots armed with guns were sent to Iraq but reportedly were never used. Nonetheless, General Robert Cone announced in 2014 that the army was considering shrinking the standard brigade combat team from 4,000 soldiers to 3,000, making up the difference with robots and drones.

So far virtually none of those robots are autonomous; a person controls each one. But the army realized this model was inefficient, so the U.S. Army Research Laboratory developed a more sophisticated robot called RoboLeader that, in the words of project chief Jessie Chen, "interprets current situations in terms of an operator's objective"—it looks, listens, senses, and determines how best to

carry out its orders—"and issues detailed command signals to a team of lower capability robots." The great advantage, as Chen explains, is that "instead of directly managing each individual robot, the human operator only deals with a single entity—RoboLeader."

Ladies and gentlemen, we have invented robotic middle management.

Robot physical skills are fast advancing on other dimensions as well. Consider a robotic hand developed by a team from Harvard, Yale, and iRobot, maker of the Roomba vacuum cleaner and many other mobile robots, including many used by the military. So fine are the robotic hand's motor skills that it can pick up a credit card from a tabletop, put a drill bit in a drill, and turn a key, all of which were previously beyond robotic abilities. "A disabled person could say to a robot with hands, 'Go to the kitchen and put my dinner in the microwave,'" one of the researchers, Harvard professor Robert Howe, told *Harvard* magazine. "Robotic hands are the real frontier, and that's where we've been pushing."

It seems that everywhere we look, computers are suddenly capable of doing things that they couldn't do and that some people thought they never would do. The less exalted skills, physical ones like folding a T-shirt, turned out to be the more challenging, but at last even they are succumbing to the combination of relentlessly increasing computing power and algorithmic skill. The number of people who wrongly believed they could never be replaced by a computer keeps growing—not slower, but faster.

THE COMPUTER KNOWS YOU'RE LYING

And yet, isn't there one last redoubt of human uniqueness, some ultimate zone of pulsing, organic personhood into which computers can never enter? Everything we've examined so far has involved abilities that originate in the left brain—logical, linear, flowchart-

able, computer-like. But what about the other side, the right side, and its specialty—emotion? It's irrational, mysterious, and we all understand it, even though we can't explain how. In addition, emotion is often the real secret sauce of success in many jobs, high-skill and low. Executives must read and respond to the emotions of customers, employees, regulators, and everyone else they deal with. A good waiter responds differently to customers who are cranky, tired, cheerful, confused, or tipsy, all without quite knowing how. Surely this is forever ours alone.

The founders of companies like Emotient and Affectiva might disagree, however. They're researchers in the field of affective computing, in which computers understand human emotion. As their work advances, our expert ability to navigate the flesh-and-blood, analog world of human feelings is looking a lot less special every day.

We express emotions in many ways—words, tone of voice, body language—but for anyone wanting to read a person's emotions continually and remotely (that is, without attaching sensors), the most useful indicator is facial expression. We aren't always speaking or moving when we feel an emotion, but we can almost always change the expression on our face. In fact, a researcher named Paul Ekman discovered decades ago that when we're feeling an emotion, it's almost impossible for us not to evince it through facial expression, even if we're trying hard to suppress it. The expression appears fleetingly, perhaps for only a fraction of a second, despite our best efforts not to betray our true feelings. Ekman became famous for his discovery of these so-called microexpressions, particularly for his work on spotting them as a way to detect lies. A TV series called *Lie to Me* was based on his work and centered on a character modeled on him.

But Ekman, one of the most cited psychologists of the twentieth century, did much more than figure out how to spot lies. He conducted the deepest analysis ever performed of all the things the human face could reveal. Your face contains about 40 muscles.

Ekman determined how each muscle could move and then figured out all the possible combinations. The total is over 10,000, meaning we can make over 10,000 facial expressions, at least in theory. In real life, about 3,000 of them have something to do with emotion. Ekman then figured out which emotion is expressed by each one of those facial muscle movement combinations. The result, which took many years to compile, was his Facial Action Coding System. For any possible combination of muscle movements in a person's face, and no matter how fleetingly the combination was visible, Ekman could tell you what that person was feeling.

Ekman built a successful business training people in law enforcement, commerce, and other fields to use his system to detect emotions. But as computer technology improved, other researchers were figuring out how to make computers do the same thing. You've noticed that even the camera in your phone can detect faces and put little boxes around them. More advanced software can examine those faces and spot the muscle movements from Ekman's system. The possibilities of such technology prompted six PhDs at the University of California at San Diego to form Emotient and to recruit Ekman to their advisory board.

Point a video camera at any person's face, and the company's Sentiment Analysis software can tell you that person's overall sentiment (positive, negative, neutral) plus display a continually updating bar chart showing levels of seven primary emotions—joy, surprise, sadness, fear, disgust, contempt, anger—and two advanced emotions, frustration and confusion (advanced because they're combinations of other emotions). Point the camera at a group of people and it analyzes all their emotions and gives you a composite readout. Incorporate the software into Google Glass, as the company has done, and the emotion readouts for anyone you're looking at appear before your eyes (and yes, several people quickly noted that the emotion you may very well detect is contempt for you because you're wearing Google Glass).

Emotient's initial target market for selling the Sentiment Analysis system was retailers, but the possibilities are obviously much broader.

Affectiva, a spin-off from MIT's Media Lab, also uses Ekman's research to analyze facial expressions, selling its software to marketers and advertisers so they can conduct consumer research online using webcams. No need to get your research subjects into a focus group and guess what they're thinking; just have them talk to you online and let their faces tell the story. A separate project within the Media Lab measures another feeling—stress in drivers—through biosensors in the steering wheel, detecting grip pressure, skin conductance, and palm sweat, and by analyzing the driver's voice as he or she speaks.

So this mysterious human ability to read another person's emotions turns out not to be all that mysterious. Computers can do it too. Technology has reached the point where doing it isn't even very hard.

That would be remarkable enough, but this supposedly unique human ability is less special still. Not only can computers detect emotion; they do it much better than we do.

Researchers led by Dr. Marian Bartlett, one of Emotient's founders, video recorded the faces of volunteers who were experiencing pain (induced by holding an arm in a bucket of ice water) or were faking the experience of pain. The videos were then shown to research subjects who were asked: Is this person really in pain or just faking it? They got it right 50 percent of the time. That is, they had no clue; they could have flipped a coin and done just as well. After the research subjects were given some training in distinguishing fake pain from real, they improved to 55 percent. But when software developed by Bartlett and her team analyzed the videos, it made the right call 85 percent of the time. The contest was not close. And nothing in this or other research suggests that pain is especially easy to fake or otherwise fundamentally unlike

other feelings. It's reasonable to infer that software would be superior at detecting other feelings as well.

Despite our deep sense that emotion is somehow uniquely ours to understand, we shouldn't be surprised that a computer can do it better. It can look at each one of those 40 muscles in the human face, and it can remember all 3,000 combinations. A video camera shoots thirty frames per second, and a computer can analyze each frame; that's just what the computer did in the detecting-fake-pain study. Even an Ekman-trained expert can occasionally miss a microexpression, but a computer never will.

Having read human emotions better than humans can, computers can then do more with that information than we can. For example, computers analyzed the faces of college students who were being trained in an uncomplicated but demanding task, a game of finding certain groups of three cards among nine cards on an iPad screen (it's harder than it sounds). The students took a test, then received training during which they were observed by humans and computers separately gauging each student's level of engagement. Then they took another test to see how much they had improved.

It turned out that the best way to predict how well the students would do on the second test was to ask the computer—based on its analysis of their facial expressions—how engaged they had been during the training. That was a better predictor than a student's performance on the first test, and it was as good as or better than the humans' judgment of engagement. And much more could be done with the computer's readings. As the researchers point out, the computer, measuring engagement moment by moment, could continually adjust the instruction to keep a student engaged and could do it individually for an unlimited number of students. Educators could mine vast troves of engagement data to learn in fine detail what grips students and what bores them. Individual

teachers, academic departments, and whole schools could monitor overall levels of student engagement.

MIT's stress-monitoring car applies the same principles in different ways. If it senses that a driver is highly stressed, it can suggest relaxing music and give the GPS guidance a more soothing tone of voice. If the driver is spacing out, it could jar him back to attention by making the steering wheel vibrate. Most intriguingly (or hilariously), the researchers even used special thermochromatic paint to make a car change its external color based on the driver's emotional state, as a signal to other drivers. Whether this technology has time to be commercialized before the arrival of autonomous vehicles makes it irrelevant is a separate question.

The power of computers to sense human emotions means, inevitably, that a machine can outdo us even in detecting our own emotions. It's surely tempting to suppose that I possess a sense of my own emotional state that no entity standing outside of me, human or electronic, could ever reach. Yet of course it isn't true. We've all had the experience of asking someone why they're in a bad mood and having that person, eyes flaming, roar, "I'm not in a bad mood!" Or more likely, we've been that person. A computer can see the obvious reality to which we're blinded. More broadly, people are notoriously awful at realizing they're stressed, depressed, angry, or in the grip of other emotions. That's why technology that warns us when we're stressed, for example, is valuable. Our supposed inner knowledge of ourselves is so frequently thwarted by inner denial that a computer is much better suited to give us a more accurate reading.

But if it's not surprising that computers can know our own emotions better than we know them ourselves, it is certainly sobering. We make life-altering decisions based on our imaginary knowledge of our own feelings. In business, love, and much else, we ask our emotions to guide us even though we're often ignorant of those

emotions. Realizing this, we sometimes ask friends or relatives for their perspective on our feelings. Life has been that way forever. What's new is that for the first time, technology can read those feeling better than people do, in themselves and in others.

The overwhelming message seems to be that, in trying to outperform technology, almost no one is safe for long. While millions of people doing a wide range of jobs are not immediately threatened, it's getting easier by the day to see how they might be. In work that is high-skill or low-skill, intensely cognitive or heavily physical, left-brain analytical or right-brain emotional, infotech is rapidly reaching a point—and in many cases has already reached a point—where it outperforms the best humans.

The conventional counterargument is that people needn't worry because the most productive use of computers is having them work with people, not replace people; humans and computers working together are even more effective than computers alone. But the cases in which this is true are always temporary; the computer alone proves superior eventually. A person working with a calculator could solve math problems infinitely faster than a calculator alone, which couldn't do a thing unless a person punched in the numbers. But today whole factories and warehouses are operated by computers alone doing complex math all by themselves; scanners and sensors supply the numbers. For a long time you and your highly computerized car could drive better than the car alone, but no more.

The favorite example to support the counterargument is chess: While IBM's Deep Blue beat world champion Garry Kasparov in 1997, by 2008 or so a person and a computer working together could beat a computer alone. You still hear this example cited as proof that people add an ineffable something that computers just can't supply. The trouble is that the argument is getting shakier by the moment. The example became famous after economist Tyler Cowen cited it in his excellent 2013 book, *Average Is Over*. But he also noted in the

book that as computers improved, the day might come when humans could no longer add any value to the computer's play, and by late 2013 he was able to post a blog entry citing evidence that the day may have arrived. At this writing, human-computer teams still beat computers some of the time, but in a domain where the world's software powerhouses like IBM and Google aren't competing. Human-computer teams can't beat computers alone at checkers, and with chess, it's only a matter of time, a few more doublings of power.

The trend holds "implications for more traditional labor markets," Cowen notes:

You might train to help a computer program read medical scans, and for thirteen years add real value with your intuition and your ability to revise the computer's mistakes or at least to get the doctor to take a closer look. But it takes more and more time for you to improve on the computer each year. And then one day . . . poof! ZMP for you.

"ZMP" means "zero marginal product"—the economists' term for when you add no value at all.

Making the trend all the more disorienting, as we keep having to remind ourselves, is that it's speeding up, not slowing down. All the technologies and developments that I've described are gee-whiz stuff as I write about them. They're probably ho-hum by the time you read this. The transition from breathtaking to boring takes only months, and that fact illustrates a crucial point. As we struggle to understand the meaning of these developments for our lives, we must force ourselves to envision larger, faster changes than the mind-blowing ones we've seen so far.

With that challenge in mind, I wanted to know what Nicholas Negroponte thought of these trends. Famous as the founder of the MIT Media Lab, he predicted decades ahead of the reality that

telephones would be transformed from wired to wireless, that TV would switch from wireless to wired, and that by now we'd all be getting our news on small digital devices. That is, he has largely gotten it right.

So I e-mailed him this query: "In five or ten years, what will people do better than computers?"

His response: "Very little, other than 'enjoy.'"

"Do we really want a world without work?"

In truth, a world that's literally without work is highly unlikely for the fundamental reason that human desires are infinite. No matter how much that technology may do for us, we will always find something to want. As Milton Friedman famously observed, you could reach a point where you pay a personal psychiatrist to follow you around. (When Henry Ford—or in some tellings, John D. Rockefeller—undertook the cruel challenge of learning golf as an adult, he reportedly paid a boy to follow him around the course and remind him to keep his head down.) We'll always find some kind of work to do. The great question as we round the fourth major turning point in the history of work is what kind of work we will find: high-value or low, the work of the psychiatrist or of the boy on the golf course.

There is a clear answer, and the strange thing about it is that to see it, we must look more closely not at computers, but at ourselves.

CHAPTER THREE


THE SURPRISING VALUE IN OUR DEEPEST NATURE

Why Being a Great Performer Is Becoming Less About What We Know and More About What We're Like.

The case in Arizona Superior Court concerned James Stone, who had been convicted twice of sexually violent crimes involving children. His prison sentence was nearing its end, and the question before the jury was whether he should be confined in a psychiatric hospital after his release—what's called civil confinement. It may seem surprising that someone could be locked up for a further, indeterminate period of time after having served his court-imposed sentence, but many states have such civil-confinement laws specifically for SVPs, or sexually violent predators or persons. The central question in a civil-confinement case is the likelihood of the person committing future acts of sexual violence, and the key testimony comes from mental health experts. A great deal is riding on how jurors respond to that testimony, because in reaching their decision they have little else to go on beyond the basic facts of the felon's crime, which may have occurred many years previously. Will these jurors release Stone into the community or lock him up in a hospital?

The court appointed a psychologist named Brent to evaluate

HUMANS ARE UNDERRATED



WHAT HIGH ACHIEVERS KNOW THAT
BRILLIANT MACHINES NEVER WILL

GEOFF COLVIN

PORTFOLIO / PENGUIN

PORTFOLIO / PENGUIN
An imprint of Penguin Random House LLC
375 Hudson Street
New York, New York 10014
penguin.com

Copyright © 2015 by Geoffrey Colvin
Penguin supports copyright. Copyright fuels creativity, encourages diverse voices,
promotes free speech, and creates a vibrant culture. Thank you for buying an
authorized edition of this book and for complying with copyright laws by
not reproducing, scanning, or distributing any part of it in any form
without permission. You are supporting writers and allowing
Penguin to continue to publish books for every reader.

ISBN: 978-1-59184-720-5

Printed in the United States of America
1 3 5 7 9 10 8 6 4 2

Set in Century MT Std
Designed by Spring Hoteling