



THE SECOND MACHINE AGE

Work, Progress, and Prosperity
in a Time of Brilliant Technologies

ERIK BRYNJOLFSSON ANDREW MCAFEE



W. W. NORTON & COMPANY

NEW YORK LONDON

To Martha Pavlakis, the love of my life.

To my parents, David McAfee and Nancy Haller, who prepared me for the second machine age by giving me every advantage a person could have.



Chapter 1 THE BIG STORIES

Chapter 2 THE SKILLS OF THE NEW MACHINES: TECHNOLOGY
RACES AHEAD

Chapter 3 MOORE'S LAW AND THE SECOND HALF OF THE
CHESSBOARD

Chapter 4 THE DIGITIZATION OF JUST ABOUT EVERYTHING

Chapter 5 INNOVATION: DECLINING OR RECOMBINING?

Chapter 6 ARTIFICIAL AND HUMAN INTELLIGENCE IN THE
SECOND MACHINE AGE

Chapter 7 COMPUTING BOUNTY

Chapter 8 BEYOND GDP

Chapter 9 THE SPREAD

Chapter 10 THE BIGGEST WINNERS: STARS AND SUPERSTARS

Chapter 11 IMPLICATIONS OF THE BOUNTY AND THE SPREAD

Chapter 12 LEARNING TO RACE *WITH* MACHINES:
RECOMMENDATIONS FOR INDIVIDUALS

Chapter 13 POLICY RECOMMENDATIONS

Chapter 14 LONG-TERM RECOMMENDATIONS

Chapter 15 TECHNOLOGY AND THE FUTURE
(Which Is Very Different from "Technology *Is* the Future")

Acknowledgments

Notes

Illustration Sources

Index



CHAPTER 2

**THE SKILLS OF
THE NEW
MACHINES:
TECHNOLOGY
RACES AHEAD**

"Any sufficiently advanced technology is indistinguishable from magic."

--Arthur C. Clarke

IN THE SUMMER OF 2012, we went for a drive in a car that had no driver.

During a research visit to Google's Silicon Valley headquarters, we got to ride in one of the company's autonomous vehicles, developed as part of its Chauffeur project. Initially we had visions of cruising in the back seat of a car that had no one in the front seat, but Google is understandably skittish about putting obviously autonomous autos on the road. Doing so might freak out pedestrians and other drivers, or attract the attention of the police. So we sat in the back while two members of the Chauffeur team rode up front.

When one of the Googlers hit the button that switched the car into fully automatic driving mode while we were headed down Highway 101, our curiosities--and self-preservation instincts--engaged. The 101 is not always a predictable or calm environment. It's nice and straight, but it's also crowded most of the time, and its traffic flows have little obvious rhyme or reason. At highway speeds the consequences of driving mistakes can be serious ones. Since we were now part of the ongoing Chauffeur experiment, these consequences were suddenly of more than just intellectual interest to us.

The car performed flawlessly. In fact, it actually provided a boring ride. It didn't speed or slalom among the other cars; it drove exactly the way we're all taught to in driver's ed. A laptop in the car provided a real-time visual representation of what the Google car 'saw' as it proceeded along the highway--all the nearby objects of which its sensors were aware. The car recognized all the surrounding vehicles, not just the nearest ones, and it remained aware of them no matter where they moved. It was a car without blind spots. But the software doing the driving was aware that cars and trucks driven by humans *do* have blind spots. The laptop screen displayed the software's best guess about where all these blind spots were and worked to stay out of them.

We were staring at the screen, paying no attention to the actual road, when traffic ahead of us came to a complete stop. The autonomous car braked smoothly in response, coming to a stop a safe distance behind the car in front, and started moving again once the rest of the traffic did. All the while the Googlers in the front seat never stopped their conversation or showed any

nervousness, or indeed much interest at all in current highway conditions. Their hundreds of hours in the car had convinced them that it could handle a little stop-and-go traffic. By the time we pulled back into the parking lot, we shared their confidence.

The New New Division of Labor

Our ride that day on the 101 was especially weird for us because, only a few years earlier, we were sure that computers would not be able to drive cars. Excellent research and analysis, conducted by colleagues who we respect a great deal, concluded that driving would remain a human task for the foreseeable future. How they reached this conclusion, and how technologies like Chauffeur started to overturn it in just a few years, offers important lessons about digital progress.

In 2004 Frank Levy and Richard Murnane published their book *The New Division of Labor*.¹ The division they focused on was between human and digital labor--in other words, between people and computers. In any sensible economic system, people should focus on the tasks and jobs where they have a comparative advantage over computers, leaving computers the work for which they are better suited. In their book Levy and Murnane offered a way to think about which tasks fell into each category.

One hundred years ago the previous paragraph wouldn't have made any sense. Back then, computers *were* humans. The word was originally a job title, not a label for a type of machine. Computers in the early twentieth century were people, usually women, who spent all day doing arithmetic and tabulating the results. Over the course of decades, innovators designed machines that could take over more and more of this work; they were first mechanical, then electro-mechanical, and eventually digital. Today, few people if any are employed simply to do arithmetic and record the results. Even in the lowest-wage countries there are no human computers, because the nonhuman ones are far cheaper, faster, and more accurate.

If you examine their inner workings, you realize that computers aren't just number crunchers, they're symbols processors. Their circuitry can be interpreted in the language of ones and zeroes, but equally validly as true or false, yes or no, or any other symbolic system. In principle, they can do all

manner of symbolic work, from math to logic to language. But digital novelists are not yet available, so people still write all the books that appear on fiction bestseller lists. We also haven't yet computerized the work of entrepreneurs, CEOs, scientists, nurses, restaurant busboys, or many other types of workers. Why not? What is it about their work that makes it harder to digitize than what human computers used to do?

Computers Are Good at Following Rules . . .

These are the questions Levy and Murnane tackled in *The New Division of Labor*, and the answers they came up with made a great deal of sense. The authors put information processing tasks--the foundation of all knowledge work--on a spectrum. At one end are tasks like arithmetic that require only the application of well-understood rules. Since computers are really good at following rules, it follows that they should do arithmetic and similar tasks.

Levy and Murnane go on to highlight other types of knowledge work that can also be expressed as rules. For example, a person's credit score is a good general predictor of whether they'll pay back their mortgage as promised, as is the amount of the mortgage relative to the person's wealth, income, and other debts. So the decision about whether or not to give someone a mortgage can be effectively boiled down to a rule.

Expressed in words, a mortgage rule might say, "If a person is requesting a mortgage of amount M and they have a credit score of V or higher, annual income greater than I or total wealth greater than W , and total debt no greater than D , then approve the request." When expressed in computer code, we call a mortgage rule like this an *algorithm*. Algorithms are simplifications; they can't and don't take everything into account (like a billionaire uncle who has included the applicant in his will and likes to rock-climb without ropes). Algorithms do, however, include the most common and important things, and they generally work quite well at tasks like predicting payback rates. Computers, therefore, can and should be used for mortgage approval.*

. . . But Lousy at Pattern Recognition

At the other end of Levy and Murnane's spectrum, however, lie information processing tasks that cannot be boiled down to rules or algorithms. According to the authors, these are tasks that draw on the human capacity for pattern recognition. Our brains are extraordinarily good at taking in information via our senses and examining it for patterns, but we're quite bad at describing or figuring out *how* we're doing it, especially when a large volume of fast-changing information arrives at a rapid pace. As the philosopher Michael Polanyi famously observed, "We know more than we can tell."² When this is the case, according to Levy and Murnane, tasks can't be computerized and will remain in the domain of human workers. The authors cite driving a vehicle in traffic as an example of such a task. As they write,

As the driver makes his left turn against traffic, he confronts a wall of images and sounds generated by oncoming cars, traffic lights, storefronts, billboards, trees, and a traffic policeman. Using his knowledge, he must estimate the size and position of each of these objects and the likelihood that they pose a hazard. . . . The truck driver [has] the schema to recognize what [he is] confronting. But articulating this knowledge and embedding it in software for all but highly structured situations are at present enormously difficult tasks. . . . Computers cannot easily substitute for humans in [jobs like driving].

So Much for *That* Distinction

We were convinced by Levy and Murnane's arguments when we read *The New Division of Labor* in 2004. We were further convinced that year by the initial results of the DARPA Grand Challenge for driverless cars.

DARPA, the Defense Advanced Research Projects Agency, was founded in 1958 (in response to the Soviet Union's launch of the *Sputnik* satellite) and tasked with spurring technological progress that might have military applications. In 2002 the agency announced its first Grand Challenge, which was to build a completely autonomous vehicle that could complete a 150-mile course through California's Mojave Desert. Fifteen entrants performed well enough in a qualifying run to compete in the main event, which was held on March 13, 2004.

The results were less than encouraging. Two vehicles didn't make it to the starting area, one flipped over *in* the starting area, and three hours into the race only four cars were still operational. The "winning" Sandstorm car from

Carnegie Mellon University covered 7.4 miles (less than 5 percent of the total) before veering off the course during a hairpin turn and getting stuck on an embankment. The contest's \$1 million prize went unclaimed, and *Popular Science* called the event "DARPA's Debacle in the Desert."³

Within a few years, however, the debacle in the desert became the 'fun on the 101' that we experienced. Google announced in an October 2010 blog post that its completely autonomous cars had for some time been driving successfully, in traffic, on American roads and highways. By the time we took our ride in the summer of 2012 the Chauffeur project had grown into a small fleet of vehicles that had collectively logged hundreds of thousands of miles with no human involvement and with only two accidents. One occurred when a person was driving the Chauffeur car; the other happened when a Google car was rear-ended (by a human driver) while stopped at a red light.⁴ To be sure, there are still many situations that Google's cars can't handle, particularly complicated city traffic or off-road driving or, for that matter, any location that has not already been meticulously mapped in advance by Google. But our experience on the highway convinced us that it's a viable approach for the large and growing set of everyday driving situations.

Self-driving cars went from being the stuff of science fiction to on-the-road reality in a few short years. Cutting-edge research explaining why they were not coming anytime soon was outpaced by cutting-edge science and engineering that brought them into existence, again in the space of a few short years. This science and engineering accelerated rapidly, going from a debacle to a triumph in a little more than half a decade.

Improvement in autonomous vehicles reminds us of Hemingway's quote about how a man goes broke: "Gradually and then suddenly."⁵ And self-driving cars are not an anomaly; they're part of a broad, fascinating pattern. Progress on some of the oldest and toughest challenges associated with computers, robots, and other digital gear was gradual for a long time. Then in the past few years it became sudden; digital gear started racing ahead, accomplishing tasks it had always been lousy at and displaying skills it was not supposed to acquire anytime soon. Let's look at a few more examples of surprising recent technological progress.

Good Listeners and Smooth Talkers

In addition to pattern recognition, Levy and Murnane highlight *complex communication* as a domain that would stay on the human side in the new division of labor. They write that, "Conversations critical to effective teaching, managing, selling, and many other occupations require the transfer and interpretation of a broad range of information. In these cases, the possibility of exchanging information with a computer, rather than another human, is a long way off."⁶

In the fall of 2011, Apple introduced the iPhone 4S featuring "Siri," an intelligent personal assistant that worked via a natural-language user interface. In other words, people talked to it just as they would talk to another human being. The software underlying Siri, which originated at the California research institute SRI International and was purchased by Apple in 2010, listened to what iPhone users were saying to it, tried to identify what they wanted, then took action and reported back to them in a synthetic voice.

After Siri had been out for about eight months, Kyle Wagner of technology blog *Gizmodo* listed some of its most useful capabilities: "You can ask about the scores of live games--'What's the score of the Giants game?'--or about individual player stats. You can also make OpenTable reservations, get Yelp scores, ask about what movies are playing at a local theater and then see a trailer. If you're busy and can't take a call, you can ask Siri to remind you to call the person back later. This is the kind of everyday task for which voice commands can actually be incredibly useful."⁷

The *Gizmodo* post ended with caution: "That actually sounds pretty cool. Just with the obvious Siri criterion: *If it actually works.*"⁸ Upon its release, a lot of people found that Apple's intelligent personal assistant didn't work well. It didn't understand what they were saying, asked for repeated clarifications, gave strange or inaccurate answers, and put them off with responses like "I'm really sorry about this, but I can't take any requests right now. Please try again in a little while." Analyst Gene Munster catalogued questions with which Siri had trouble:

* ***Where is Elvis buried?*** Responded, "I can't answer that for

you." It thought the person's name was Elvis Buried.

- * ***When did the movie Cinderella come out?*** Responded with a movie theater search on Yelp.
- * ***When is the next Halley's Comet?*** Responded, "You have no meetings matching Halley's."
- * ***I want to go to Lake Superior.*** Responded with directions to the company Lake Superior X-Ray.⁹

Siri's sometimes bizarre and frustrating responses became well known, but the power of the technology is undeniable. It can come to your aid exactly when you need it. On the same trip that afforded us some time in an autonomous car, we saw this firsthand. After a meeting in San Francisco, we hopped in our rental car to drive down to Google's headquarters in Mountain View. We had a portable GPS device with us, but didn't plug it in and turn it on because we thought we knew how to get to our next destination.

We didn't, of course. Confronted with an Escherian maze of elevated highways, off-ramps, and surface streets, we drove around looking for an on-ramp while tensions mounted. Just when our meeting at Google, this book project, and our professional relationship seemed in serious jeopardy, Erik pulled out his phone and asked Siri for "directions to U.S. 101 South." The phone responded instantly and flawlessly: the screen turned into a map showing where we were and how to find the elusive on-ramp.

We could have pulled over, found the portable GPS and turned it on, typed in our destination, and waited for our routing, but we didn't want to exchange information that way. We wanted to speak a question and hear and see (because a map was involved) a reply. Siri provided exactly the natural language interaction we were looking for. A 2004 review of the previous half-century's research in automatic speech recognition (a critical part of natural language processing) opened with the admission that "Human-level speech recognition has proved to be an elusive goal," but less than a decade later major elements of that goal have been reached. Apple and other companies have made robust natural language processing technology available to hundreds of millions of people via their mobile phones.¹⁰ As noted by Tom Mitchell, who heads the machine-learning department at Carnegie Mellon University: "We're at the beginning of a ten-year period

where we're going to transition from computers that can't understand language to a point where computers can understand quite a bit about language."¹¹

Digital Fluency: The Babel Fish Goes to Work

Natural language processing software is still far from perfect, and computers are not yet as good as people at complex communication, but they're getting better all the time. And in tasks like translation from one language to another, surprising developments are underway: while computers' communication abilities are not as deep as those of the average human being, they're much broader.

A person who speaks more than one language can usually translate between them with reasonable accuracy. Automatic translation services, on the other hand, are impressive but rarely error-free. Even if your French is rusty, you can probably do better than Google Translate with the sentence "Monty Python's 'Dirty Hungarian Phrasebook' sketch is one of their funniest ones." Google offered, "Sketch des Monty Python 'Phrasebook sale hongrois' est l'un des plus drols les leurs." This conveys the main gist, but has serious grammatical problems.

There is less chance you could have made progress translating this sentence (or any other) into Hungarian, Arabic, Chinese, Russian, Norwegian, Malay, Yiddish, Swahili, Esperanto, or any of the other sixty-three languages besides French that are part of the Google Translate service. But Google will attempt a translation of text from any of these languages into any other, instantaneously and at no cost for anyone with Web access.¹² The Translate service's smartphone app lets users speak more than fifteen of these languages into the phone and, in response, will produce synthesized, translated speech in more than half of the fifteen. It's a safe bet that even the world's most multilingual person can't match this breadth.

For years instantaneous translation utilities have been the stuff of science fiction (most notably *The Hitchhiker's Guide to the Galaxy's* Babel Fish, a strange creature that once inserted in the ear allows a person to understand speech in any language).¹³ Google Translate and similar services are making it a reality today. In fact, at least one such service is being used right now to

facilitate international customer service interactions. The translation services company Lionbridge has partnered with IBM to offer GeoFluent, an online application that instantly translates chats between customers and troubleshooters who do not share a language. In an initial trial, approximately 90 percent of GeoFluent users reported that it was good enough for business purposes.¹⁴

Human Superiority in *Jeopardy!*

Computers are now combining pattern matching with complex communication to quite literally beat people at their own games. In 2011, the February 14 and 15 episodes of the TV game show *Jeopardy!* included a contestant that was not a human being. It was a supercomputer called Watson, developed by IBM specifically to play the game (and named in honor of legendary IBM CEO Thomas Watson, Sr.). *Jeopardy!* debuted in 1964 and in 2012 was the fifth most popular syndicated TV program in America.¹⁵ On a typical day almost 7 million people watch host Alex Trebek ask trivia questions on various topics as contestants vie to be the first to answer them correctly.*

The show's longevity and popularity stem from its being easy to understand yet extremely hard to play well. Almost everyone knows the answers to some of the questions in a given episode, but very few people know the answers to almost all of them. Questions cover a wide range of topics, and contestants are not told in advance what those topics will be. Players also have to be simultaneously fast, bold, and accurate--fast because they compete against one another for the chance to answer each question; bold because they have to try to answer a lot of questions, especially harder ones, in order to accumulate enough money to win; and accurate because money is subtracted for each incorrect answer.

Jeopardy!'s producers further challenge contestants with puns, rhymes, and other kinds of wordplay. A clue might ask, for example, for "A rhyming reminder of the past in the city of the NBA's Kings."¹⁶ To answer correctly, a player would have to know what the acronym NBA stood for (in this case, it's the National Basketball Association, not the National Bank Act or chemical compound n-Butylamine), which city the NBA's Kings play in (Sacramento),

and that the clue's demand for a *rhyming* reminder of the past meant that the right answer is "What is a Sacramento memento?" instead of a "Sacramento souvenir" or any other factually correct response. Responding correctly to clues like these requires mastery of pattern matching and complex communication. And winning at *Jeopardy!* requires doing both things repeatedly, accurately, and almost instantaneously.

During the 2011 shows, Watson competed against Ken Jennings and Brad Rutter, two of the best knowledge workers in this esoteric industry. Jennings won *Jeopardy!* a record seventy-four times in a row in 2004, taking home more than \$3,170,000 in prize money and becoming something of a folk hero along the way.¹⁷ In fact, Jennings is sometimes given credit for the existence of Watson.¹⁸ According to one story circulating within IBM, Charles Lickel, a research manager at the company interested in pushing the frontiers of artificial intelligence, was having dinner in a steakhouse in Fishkill, New York, one night in the fall of 2004. At 7 p.m., he noticed that many of his fellow diners got up and went into the adjacent bar. When he followed them to find out what was going on, he saw that they were clustered in front of the bar's TV watching Jennings extend his winning streak beyond fifty matches. Lickel saw that a match between Jennings and a *Jeopardy!*-playing supercomputer would be extremely popular, in addition to being a stern test of a computer's pattern matching and complex communication abilities.

Since *Jeopardy!* is a three-way contest, the ideal third contestant would be Brad Rutter, who beat Jennings in the show's 2005 Ultimate Tournament of Champions and won more than \$3,400,000.¹⁹ Both men had packed their brains with information of all kinds, were deeply familiar with the game and all of its idiosyncrasies, and knew how to handle pressure.

These two humans would be tough for any machine to beat, and the first versions of Watson weren't even close. Watson could be 'tuned' by its programmers to be either more aggressive in answering questions (and hence more likely to be wrong) or more conservative and accurate. In December 2006, shortly after the project started, when Watson was tuned to try to answer 70 percent of the time (a relatively aggressive approach) it was only able to come up with the right response approximately 15 percent of the time. Jennings, in sharp contrast, answered about 90 percent of questions correctly in games when he buzzed in first (in other words, won the right to respond)

70 percent of the time.²⁰

But Watson turned out to be a very quick learner. The supercomputer's performance on the aggression vs. accuracy tradeoff improved quickly, and by November 2010, when it was aggressive enough to win the right to answer 70 percent of a simulated match's total questions, it answered about 85 percent of them correctly. This was impressive improvement, but it still didn't put the computer in the same league as the best human players. The Watson team kept working until mid-January of 2011, when the matches were recorded for broadcast in February, but no one knew how well their creation would do against Jennings and Rutter.

Watson trounced them both. It correctly answered questions on topics ranging from "Olympic Oddities" (responding "pentathlon" to "A 1976 entry in the 'modern' this was kicked out for wiring his epee to score points without touching his foe") to "Church and State" (realizing that the answers all contained one or the other of these words, the computer answered "gestate" when told "It can mean to develop gradually in the mind or to carry during pregnancy"). While the supercomputer was not perfect (for example, it answered "chic" instead of "class" when asked about "stylish elegance, or students who all graduated in the same year" as part of the category "Alternate Meanings"), it was very good.

Watson was also extremely fast, repeatedly buzzing in before Jennings and Rutter to win the right to answer questions. In the first of the two games played, for example, Watson buzzed in first 43 times, then answered correctly 38 times. Jennings and Rutter *combined* to buzz in only 33 times over the course of the same game.²¹

At the end of the two-day tournament, Watson had amassed \$77,147, more than three times as much as either of its human opponents. Jennings, who came in second, added a personal note on his answer to the tournament's final question: "I for one welcome our new computer overlords." He later elaborated, "Just as factory jobs were eliminated in the twentieth century by new assembly-line robots, Brad and I were the first knowledge-industry workers put out of work by the new generation of 'thinking' machines. 'Quiz show contestant' may be the first job made redundant by Watson, but I'm sure it won't be the last."²²

The Paradox of Robotic 'Progress'

A final important area where we see a rapid recent acceleration in digital improvement is robotics--building machines that can navigate through and interact with the physical world of factories, warehouses, battlefields, and offices. Here again we see progress that was very gradual, then sudden.

The word *robot* entered the English language via the 1921 Czech play, *R.U.R.* (Rossum's "Universal" Robots) by Karel Capek, and automatons have been an object of human fascination ever since.²³ During the Great Depression, magazine and newspaper stories speculated that robots would wage war, commit crimes, displace workers, and even beat boxer Jack Dempsey.²⁴ Isaac Asimov coined the term *robotics* in 1941 and provided ground rules for the young discipline the following year with his famous Three Laws of Robotics:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.²⁵

Asimov's enormous influence on both science fiction and real-world robot-making has persisted for seventy years. But one of those two communities has raced far ahead of the other. Science fiction has given us the chatty and loyal R2-D2 and C-3PO, *Battlestar Galactica's* ominous Cylons, the terrible Terminator, and endless varieties of androids, cyborgs, and replicants. Decades of robotics research, in contrast, gave us Honda's ASIMO, a humanoid robot best known for a spectacularly failed demo that showcased its inability to follow Asimov's third law. At a 2006 presentation to a live audience in Tokyo, ASIMO attempted to walk up a shallow flight of stairs that had been placed on the stage. On the third step, the robot's knees buckled and it fell over backward, smashing its faceplate on the floor.²⁶

ASIMO has since recovered and demonstrated skills like walking up and down stairs, kicking a soccer ball, and dancing, but its shortcomings highlight a broad truth: a lot of the things humans find easy and natural to do in the

physical world have been remarkably difficult for robots to master. As the roboticist Hans Moravec has observed, "It is comparatively easy to make computers exhibit adult-level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility."²⁷

This situation has come to be known as Moravec's paradox, nicely summarized by Wikipedia as "the discovery by artificial intelligence and robotics researchers that, contrary to traditional assumptions, high-level reasoning requires very little computation, but low-level sensorimotor skills require enormous computational resources."^{28*} Moravec's insight is broadly accurate, and important. As the cognitive scientist Steven Pinker puts it, "The main lesson of thirty-five years of AI research is that the hard problems are easy and the easy problems are hard. . . . As the new generation of intelligent devices appears, it will be the stock analysts and petrochemical engineers and parole board members who are in danger of being replaced by machines. The gardeners, receptionists, and cooks are secure in their jobs for decades to come."²⁹

Pinker's point is that robotics experts have found it fiendishly difficult to build machines that match the skills of even the least-trained manual worker. iRobot's Roomba, for example, can't do everything a maid does; it just vacuums the floor. More than ten million Roombas have been sold, but none of them is going to straighten the magazines on a coffee table.

When it comes to work in the physical world, humans also have a huge flexibility advantage over machines. Automating a single activity, like soldering a wire onto a circuit board or fastening two parts together with screws, is pretty easy, but that task must remain constant over time and take place in a 'regular' environment. For example, the circuit board must show up in exactly the same orientation every time. Companies buy specialized machines for tasks like these, have their engineers program and test them, then add them to their assembly lines. Each time the task changes--each time the location of the screw holes move, for example--production must stop until the machinery is reprogrammed. Today's factories, especially large ones in high-wage countries, are highly automated, but they're not full of general-purpose robots. They're full of dedicated, specialized machinery that's expensive to buy, configure, and reconfigure.

Rethinking Factory Automation

Rodney Brooks, who co-founded iRobot, noticed something else about modern, highly automated factory floors: people are scarce, but they're not absent. And a lot of the work they do is repetitive and mindless. On a line that fills up jelly jars, for example, machines squirt a precise amount of jelly into each jar, screw on the top, and stick on the label, but a person places the empty jars on the conveyor belt to start the process. Why hasn't this step been automated? Because in this case the jars are delivered to the line twelve at a time in cardboard boxes that don't hold them firmly in place. This imprecision presents no problem to a person (who simply sees the jars in the box, grabs them, and puts them on the conveyor belt), but traditional industrial automation has great difficulty with jelly jars that don't show up in exactly the same place every time.

In 2008 Brooks founded a new company, Rethink Robotics, to pursue and build *un*traditional industrial automation: robots that can pick and place jelly jars and handle the countless other imprecise tasks currently done by people in today's factories. His ambition is to make some progress against Moravec's paradox. What's more, Brooks envisions creating robots that won't need to be programmed by high-paid engineers; instead, the machines can be taught to do a task (or retrained to do a new one) by shop floor workers, each of whom need less than an hour of training to learn how to instruct their new mechanical colleagues. Brooks's machines are cheap, too. At about \$20,000, they're a small fraction of the cost of current industrial robots. We got a sneak peek at these potential paradox-busters shortly before Rethink's public unveiling of their first line of robots, named Baxter. Brooks invited us to the company's headquarters in Boston to see these automatons, and to see what they could do.

Baxter is instantly recognizable as a humanoid robot. It has two burly, jointed arms with claw-like grips for hands; a torso; and a head with an LCD face that swivels to 'look at' the nearest person. It doesn't have legs, though; Rethink sidestepped the enormous challenges of automatic locomotion by putting Baxter on wheels and having it rely on people to get from place to place. The company's analyses suggest that it can still do lots of useful work

without the ability to move under his own power.

To train Baxter, you grab it by the wrist and guide the arm through the motions you want it to carry out. As you do this, the arm seems weightless; its motors are working so you don't have to. The robot also maintains safety; the two arms can't collide (the motors resist you if you try to make this happen) and they automatically slow down if Baxter senses a person within their range. These and many other design features make working with this automaton a natural, intuitive, and nonthreatening experience. When we first approached it, we were nervous about catching a robot arm to the face, but this apprehension faded quickly, replaced by curiosity.

Brooks showed us several Baxters at work in the company's demo area. They were blowing past Moravec's paradox--sensing and manipulating lots of different objects with 'hands' ranging from grips to suction cups. The robots aren't as fast or fluid as a well-trained human worker at full speed, but they might not need to be. Most conveyor belts and assembly lines do not operate at full human speed; they would tire people out if they did.

Baxter has a few obvious advantages over human workers. It can work all day every day without needing sleep, lunch, or coffee breaks. It also won't demand healthcare from its employer or add to the payroll tax burden. And it can do two completely unrelated things at once; its two arms are capable of operating independently.

Coming Soon to Assembly Lines, Warehouses, and Hallways Near You

After visiting Rethink and seeing Baxter in action, we understood why Texas Instruments Vice President Remi El-Ouazzane said in early 2012, "We have a firm belief that the robotics market is on the cusp of exploding." There's a lot of evidence to support his view. The volume and variety of robots in use at companies is expanding rapidly, and innovators and entrepreneurs have recently made deep inroads against Moravec's paradox.³⁰

Kiva, another young Boston-area company, has taught its automatons to move around warehouses safely, quickly, and effectively. Kiva robots look like metal ottomans or squashed R2-D2s. They scuttle around buildings at

about knee-height, staying out of the way of humans and one another. They're low to the ground so they can scoot underneath shelving units, lift them up, and bring them to human workers. After these workers grab the products they need, the robot whisks the shelf away and another shelf-bearing robot takes its place. Software tracks where all the products, shelves, robots, and people are in the warehouse, and orchestrates the continuous dance of the Kiva automatons. In March of 2012, Kiva was acquired by Amazon--a leader in advanced warehouse logistics--for more than \$750 million in cash.³¹

Boston Dynamics, yet another New England startup, has tackled Moravec's paradox head-on. The company builds robots aimed at supporting American troops in the field by, among other things, carrying heavy loads over rough terrain. Its BigDog, which looks like a giant metal mastiff with long skinny legs, can go up steep hills, recover from slips on ice, and do other very dog-like things. Balancing a heavy load on four points while moving over an uneven landscape is a truly nasty engineering problem, but Boston Dynamics has been making good progress.

As a final example of recent robotic progress, consider the Double, which is about as different from the BigDog as possible. Instead of trotting through rough enemy terrain, the Double rolls over cubicle carpets and hospital hallways carrying an iPad. It's essentially an upside-down pendulum with motorized wheels at the bottom and a tablet at the top of a four-to five-foot stick. The Double provides telepresence--it lets the operator 'walk around' a distant building and see and hear what's going on. The camera, microphone, and screen of the iPad serve as the eyes, ears, and face of the operator, who sees and hears what the iPad sees and hears. The Double itself acts as the legs, transporting the whole assembly around in response to commands from the operator. Double Robotics calls it "the simplest, most elegant way to be somewhere else in the world without flying there." The first batch of Doubles, priced at \$2,499, sold out soon after the technology was announced in the fall of 2012.³²

The next round of robotic innovation might put the biggest dent in Moravec's paradox ever. In 2012 DARPA announced another Grand Challenge; instead of autonomous cars, this one was about automatons. The DARPA Robotics Challenge (DRC) combined tool use, mobility, sensing, telepresence, and many other long-standing challenges in the field. According

to the website of the agency's Tactical Technology Office,

The primary technical goal of the DRC is to develop ground robots capable of executing complex tasks in dangerous, degraded, human-engineered environments. Competitors in the DRC are expected to focus on robots that can use standard tools and equipment commonly available in human environments, ranging from hand tools to vehicles, with an emphasis on adaptability to tools with diverse specifications.³³

With the DRC, DARPA is asking the robotics community to build and demonstrate high-functioning humanoid robots by the end of 2014. According to an initial specification supplied by the agency, they will have to be able to drive a utility vehicle, remove debris blocking an entryway, climb a ladder, close a valve, and replace a pump.³⁴ These seem like impossible requirements, but we've been assured by highly knowledgeable colleagues--ones competing in the DRC, in fact--that they'll be met. Many saw the 2004 Grand Challenge as instrumental in accelerating progress with autonomous vehicles. There's an excellent chance that the DRC will be similarly important at getting us past Moravec's paradox.

More Evidence That We're at an Inflection Point

Self-driving cars, *Jeopardy!* champion supercomputers, and a variety of useful robots have all appeared just in the past few years. And these innovations are not just lab demos; they're showing off their skills and abilities in the messy real world. They contribute to the impression that we're at an inflection point--a bend in the curve where many technologies that used to be found only in science fiction are becoming everyday reality. As many other examples show, this is an accurate impression.

On the *Star Trek* television series, devices called tricorders were used to scan and record three kinds of data: geological, meteorological, and medical. Today's consumer smartphones serve all these purposes; they can be put to work as seismographs, real-time weather radar maps, and heart-and-breathing-rate monitors.³⁵ And, of course, they're not limited to these domains. They also work as media players, game platforms, reference works, cameras, and GPS devices. On *Star Trek*, tricorders and person-to-person communicators were separate devices, but in the real world the two have

merged in the smartphone. They enable their users to simultaneously access and generate huge amounts of information as they move around. This opens up the opportunity for innovations that venture capitalist John Doerr calls "SoLoMo"--social, local, and mobile.³⁶

Computers historically have been very bad at writing real prose. In recent times they have been able to generate grammatically correct but meaningless sentences, a state of affairs that's been mercilessly exploited by pranksters. In 2008, for example, the International Conference on Computer Science and Software Engineering accepted the paper "Towards the Simulation of E-commerce" and invited its author to chair a session. This paper was 'written' by SCIGen, a program from the MIT Computer Science and Artificial Intelligence Lab that "generates random Computer Science research papers." SCIGen's authors wrote that, "Our aim here is to maximize amusement, rather than coherence," and after reading the abstract of "Towards the Simulation of E-commerce" it's hard to argue with them.³⁷

Recent advances in cooperative technology and classical communication are based entirely on the assumption that the Internet and active networks are not in conflict with object-oriented languages. In fact, few information theorists would disagree with the visualization of DHTs that made refining and possibly simulating 8 bitarchitectures a reality, which embodies the compelling principles of electrical engineering.³⁸

Recent developments make clear, though, that not all computer-generated prose is nonsensical. Forbes.com has contracted with the company Narrative Science to write the corporate earnings previews that appear on the website. These stories are all generated by algorithms without human involvement. And they're indistinguishable from what a human would write:

Forbes Earning Preview: H.J. Heinz

A quality first quarter earnings announcement could push shares of H.J. Heinz (HNZ) to a new 52-week high as the price is just 49 cents off the milestone heading into the company's earnings release on Wednesday, August 29, 2012.

The Wall Street consensus is 80 cents per share, up 2.6 percent from a year ago when H.J reported earnings of 78 cents per share.

The consensus estimate remains unchanged over the past month, but it has decreased from three months ago when it was 82 cents. Analysts are expecting earnings of \$3.52 per share for the fiscal year. Analysts project

revenue to fall 0.3 percent year-over-year to \$2.84 billion for the quarter, after being \$2.85 billion a year ago. For the year, revenue is projected to roll in at \$11.82 billion.³⁹

Even computer peripherals like printers are getting in on the act, demonstrating useful capabilities that seem straight out of science fiction. Instead of just putting ink on paper, they are making complicated three-dimensional parts out of plastic, metal, and other materials. 3D printing, also sometimes called "additive manufacturing," takes advantage of the way computer printers work: they deposit a very thin layer of material (ink, traditionally) on a base (paper) in a pattern determined by the computer.

Innovators reasoned that there is nothing stopping printers from depositing layers one on top of the other. And instead of ink, printers can also deposit materials like liquid plastic that gets cured into a solid by ultraviolet light. Each layer is very thin--somewhere around one-tenth of a millimeter--but over time a three-dimensional object takes shape. And because of the way it is built up, this shape can be quite complicated--it can have voids and tunnels in it, and even parts that move independently of one another. At the San Francisco headquarters of Autodesk, a leading design software company, we handled a working adjustable wrench that was printed as a single part, no assembly required.⁴⁰

This wrench was a demonstration product made out of plastic, but 3D printing has expanded into metals as well. Autodesk CEO Carl Bass is part of the large and growing community of additive manufacturing hobbyists and tinkerers. During our tour of his company's gallery, a showcase of all the products and projects enabled by Autodesk software, he showed us a beautiful metal bowl he designed on a computer and had printed out. The bowl had an elaborate lattice pattern on its sides. Bass said that he'd asked friends of his who were experienced in working with metal--sculptors, ironworkers, welders, and so on--how the bowl was made. None of them could figure out how the lattice was produced. The answer was that a laser had built up each layer by fusing powdered metal.

3D printing today is not just for art projects like Bass's bowl. It's used by countless companies every day to make prototypes and model parts. It's also being used for final parts ranging from plastic vents and housings on NASA's next-generation Moon rover to a metal prosthetic jawbone for an eighty-

three-year-old woman. In the near future, it might be used to print out replacement parts for faulty engines on the spot instead of maintaining stockpiles of them in inventory. Demonstration projects have even shown that the technique could be used to build concrete houses.⁴¹

Most of the innovations described in this chapter have occurred in just the past few years. They've taken place in areas where improvement had been frustratingly slow for a long time, and where the best thinking often led to the conclusion that it wouldn't speed up. But then digital progress became sudden after being gradual for so long. This happened in multiple areas, from artificial intelligence to self-driving cars to robotics.

How did this happen? Was it a fluke--a confluence of a number of lucky one-time advances? No, it was not. The digital progress we've seen recently is certainly impressive, but it's just a small indication of what's to come. It's the dawn of the second machine age. To understand why it's unfolding now, we need to understand the nature of technological progress in the era of digital hardware, software, and networks. In particular, we need to understand its three key characteristics: that it is *exponential*, *digital*, and *combinatorial*. The next three chapters will discuss each of these in turn.

* In the years leading up to the Great Recession that began in 2007, companies were giving mortgages to people with lower and lower credit scores, income, and wealth, and higher and higher debt levels. In other words, they either rewrote or ignored their previous mortgage approval algorithms. It wasn't that the old mortgage algorithms stopped working; it was that they stopped being used.

* To be precise, Trebek reads answers and the contestants have to state the question that would give rise to this answer.

* Sensorimotor skills are those that involve sensing the physical world and controlling the body to move through it.