"The Rise of AI" The End of Code
Wired Magazine June 2016

tus to those who speak it. They have access to what in a more mechanical age would have been called the levers of power. "If you control the code, you control the world," wrote futurist Marc Goodman. (In *Bloomberg Businessweek*, Paul Ford was slightly more circumspect: "If coders don't run the world, they run the things that run the world." Tomato, tomahto.)

But whether you like this state of affairs or hate it—whether you're a member of the coding elite or someone who barely feels competent to futz with the settings on your phone—don't get used to it. Our machines are starting to speak a different language now, one that even the best coders can't fully understand.

Over the past several years, the biggest tech companies in Silicon Valley have aggressively pursued an approach to computing called machine learning. In traditional programming, an engineer writes explicit, step-by-step instructions for the computer to follow. With machine learning, programmers don't encode computers with instructions. They *train* them. If you want to teach a neural network to recognize a cat, for instance, you don't tell it to look for whiskers, ears, fur, and eyes. You simply show it thousands and thousands of photos of cats, and eventually it works things out. If it keeps misclassifying foxes as cats, you don't rewrite the code. You just keep coaching it.

This approach is not new—it's been around for decades—but it has recently become immensely more powerful, thanks in part to the rise of deep neural networks, massively distributed computational systems that mimic the multilayered connections of neurons in the brain. And already, whether you realize it or not, machine learning powers large swaths of our online activity. Facebook uses it to determine which stories show up in your News Feed, and Google Photos uses it to identify faces. Machine learning runs Microsoft's Skype Translator, which converts speech to different languages in real time. Self-driving cars use machine learning to avoid accidents. Even Google's search engine—for so many years a towering edifice of human-written rules—has begun to rely on these deep neural networks. In February the company replaced its longtime head of search with machine-learning expert John Giannandrea, and it has initiated a major program to retrain its engineers in these new techniques. "By building learning systems," Giannandrea told reporters this fall, "we don't have to write these rules anymore."

But here's the thing: With machine learning, the engineer never knows precisely how the computer accomplishes its tasks. The neural network's operations are largely opaque and inscrutable. It is, in other words, a black box. And as these black boxes assume responsibility for more and more of our daily digital tasks, they are not only going to change our relationship to technology—they are going to change how we think about ourselves, our world, and our place within it.

If in the old view programmers were like gods, authoring the laws that govern computer systems, now they're like parents or dog trainers. And as any parent or dog owner can tell you, that is a much more mysterious relationship to find yourself in.

Andy Rubin is an inveterate tinkerer and coder. The cocreator of the Android operating system, Rubin is notorious in Silicon Valley for filling his workplaces and home with robots. He programs them himself. "I got into computer science when I was very young, and I loved it because I could disappear in the world of the computer. It was a clean slate, a blank canvas, and I could create something from scratch," he says. "It gave me full control of a world that I played in for many, many years."

Now, he says, that world is coming to an end. Rubin is excited about the rise of machine learning—his new company, Playground Global, invests in machine-learning startups and is positioning itself to lead the spread of intelligent devices—but it saddens him a little too. Because machine learning changes what it means to be an engineer.

"People don't linearly write the programs," Rubin says. "After a neural network learns how to do speech recognition, a programmer can't go in and look at it and see how that happened. It's just like your brain. You can't cut your head off and see what you're thinking." When engineers do peer into a deep neural network, what they see is an ocean of math: a massive, multilayer set of calculus problems that—by constantly deriving the relationship between billions of data points—generate guesses about the world.

Artificial intelligence wasn't supposed to work this way. Until a few years ago, mainstream AI researchers assumed that to create intelligence, we just had to imbue a machine with the right logic. Write enough rules and eventually we'd create a system sophisticated enough to understand the world. They largely ignored, even vilified, early propo-

0 7 7

The Rise of AI

Could it fight cancer?

Machines are learning to read your genetic future.

In 2002, as scientists were wrapping up the first complete sequence of the human genome, Brendan Frey found out his unborn child had a genetic anomaly. "I learned it could be nothing or it could be a huge problem," he says. "Basically it was an informational wasteland." At the time, Frey was a machine-learning researcher at the University of Toronto, working on artificial intelligence software that could interpret speech and visually identify objects. What scientists really needed, he realized, was superhuman intelligence that could *analyze* genomes. So in 2014, Frey cofounded Deep Genomics with the aim of applying machine learning to DNA.

Even today, scientists are baffled by why a gene might cause breast cancer in one person but not the next. Part of the answer likely lies in the 98 percent of the genome that doesn't code for genes—once called junk DNA—that somehow influences whether genes nearby or even far away are turned on. Deep Genomics uses machine learning to probe those interactions and figure out how they translate into genetic destiny.

In his previous research, Frey built AI systems that used probabilistic calculations and huge amounts of data to emulate what any human does when they read a word or recognize a face. Now he's using the same kind of approach to build a system that can emulate what a cell does when it reads a genome and generates a new molecule. That's the first challenge, anyway. Next: cutting cancer and other diseases off at the pass.

—SARAH ZHANG

illustrations by Edward C. Monaghan

nents of machine learning, who argued in favor of plying machines with data until they reached their own conclusions. For years computers weren't powerful enough to really prove the merits of either approach, so the argument became a philosophical one. "Most of these debates were based on fixed beliefs about how the world had to be organized and how the brain worked," says Sebastian Thrun, the former Stanford AI professor who created Google's self-driving car. "Neural nets had no symbols or rules, just numbers. That alienated a lot of people."

The implications of an unparsable machine language aren't just philosophical. For the past two decades, learning to code has been one of the surest routes to reliable employment—a fact not lost on all those parents enrolling their kids in after-school code academies. But a world run by neurally networked deep-learning machines requires a different workforce. Analysts have already started worrying about the impact of AI on the job market, as machines render old skills irrelevant. Programmers might soon get a taste of what that feels like themselves.

"I was just having a conversation about that this morning," says tech guru Tim O'Reilly when I ask him about this shift. "I was pointing out how different programming jobs would be by the time all these STEM-educated kids grow up." Traditional coding won't disappear completely—indeed, O'Reilly predicts that we'll still need coders for a long time yet—but there will likely be less of it, and it will become a meta skill, a way of creating what Oren Etzioni, CEO of the Allen Institute for Artificial Intelligence, calls the "scaffolding" within which machine learning can operate. Just as Newtonian physics wasn't obviated by the discovery of quantum mechanics, code will remain a powerful, if incomplete, tool set to explore the world. But when it comes to powering specific functions, machine learning will do the bulk of the work for us.

Of course, humans still have to train these systems. But for now, at least, that's a rarefied skill. The job requires both a high-level grasp of mathematics and an intuition for pedagogical give-and-take. "It's almost like an art form to get the best out of these systems," says Demis Hassabis, who leads Google's DeepMind AI team. "There's only a few hundred people in the world that can do that really well." But even that tiny number has been enough to transform the tech industry in just a couple of years.

Whatever the professional implications of this shift, the cultural consequences will be even bigger. If the rise of human-written software led to the cult of the engineer, and to the notion that human experience can ultimately be reduced to a series of comprehensible instructions, machine learning kicks the pendulum in the opposite direction. The code that runs the universe may defy human analysis. Right now Google, for example, is facing an antitrust investigation in Europe that accuses the company of exerting undue influence over its search results. Such a charge will be difficult to prove when even the company's own engineers can't say exactly how its search algorithms work in the first place.

This explosion of indeterminacy has been a long time coming. It's not news that even simple algorithms can create unpredictable emergent behavior—an insight that goes back to chaos theory and random number generators. Over the past few years, as networks have grown more intertwined and their functions more complex, code has come to seem more like an alien force, the ghosts in the machine ever more elusive and ungovernable. Planes grounded for no reason. Seemingly unpreventable flash crashes in the stock market. Rolling blackouts.

These forces have led technologist Danny Hillis to declare the end of the age of Enlightenment, our centuries-long faith in logic, determinism, and control over nature. Hillis says we're shifting to what he calls the age of Entanglement. "As our technological and institutional creations have become more complex, our relationship to them has changed," he wrote in the *Journal of Design and Science*. "Instead of being masters of our creations, we have learned to bargain with them, cajoling and guiding them in the general direction of our goals. We have built our own jungle, and it has a life of its own." The rise of machine learning is the latest—and perhaps the last—step in this journey.

This can all be pretty frightening. After all, coding was at least the kind of thing that a regular person could imagine picking up at a boot camp. Coders were at least *human*. Now the technological elite is even smaller, and their command over their creations has waned and become indirect. Already the companies that build this stuff find it behaving in ways that are hard to govern. Last summer, Google rushed to apologize when its photo recognition engine started tagging images of black people as gorillas. The company's blunt first fix was to keep the system from labeling *anything* as a gorilla.

The Rise of AI

Will it destroy us?

Open source may help computers not be evil.

Who will defend humanity against an evil artificial intelligence that wants to rule the world? Elon Musk. Obviously. With venture capitalist Sam Altman, the Tesla CEO has built a billion-dollar organization to fight malicious AI. Their secret weapon: *more* AI.

Wait. What? Yup: The group, OpenAI, is building AI software and giving it away. The idea is that putting more AI out in the world—and allowing everyone the freedom to tweak it—will mean no company or government will have a monopoly. An AI could still go rogue, says Greg Brockman, OpenAI's chief technology officer, "but if there are many agents with about the same capabilities, they could keep any one bad actor in check."

Before the robopocalypse, a bigger AI tent could have other benefits. Companies and individuals could find new, creative ways to use it, and a wider range of backgrounds might help make AIs that benefit the whole world. "What we're actually doing when we code is describing our world from our particular perspective," says Damien Williams, a Kennesaw State University philosopher specializing in the ethics of nonhuman consciousness. "Whatever assumptions and biases we have in ourselves are very likely to be replicated in that code." Remember how users gamed Microsoft's chatbot Tay into posting racist tweets? Human prejudice can warp artificial minds.

Still, the open source world is hardly diverse; OpenAI will have to work to be inclusive. It takes a village to raise a robot.

—KLINT FINLEY

even elated, that the machine has achieved what so many thought it couldn't. But even he feels his humanness rise. He starts to hope that Lee will win one.

Two hours into game four, Lee is deep in another hole. He plays an aggressive game, attacking particular areas of the sprawling game board. But AlphaGo plays a more expansive style, taking a more holistic approach that weighs the entire board. In Move 37, AlphaGo placed its black stone in an area near only one other stone, away from the main action. Once again, in game four, the machine is using this enigmatic approach to take control of the contest.

AlphaGo has already won the tournament. Lee isn't playing for the win anymore. He's playing for humanity. Seventy-seven moves in, he seems to stall. He rests his chin in his right hand. He sways forward and back. He swivels in his chair and rubs the back of his neck. Two minutes pass, then four, then six.

Then, still gripping the back of his neck with his left hand, he strikes. With the first two fingers of his right hand, Lee puts a white stone near the very center of the board, directly between two black stones. It's the 78th stone on the board, a "wedge move" between two vast and crowded swaths of territory. It effectively cuts AlphaGo's defenses in half. And the machine blinks. Not literally, of course. But its next move is horrendous. Lee shoots a pointed stare at Huang, as if Huang is the opponent rather than a billion circuits.

In AlphaGo's control room, the people running the machine stop what they're doing and stare at their monitors. Before Lee's brilliant Move 78, AlphaGo was putting its chances of winning at 70 percent. Eight moves later, the odds drop off the table. Suddenly AlphaGo isn't Deep Blue's successor—it's Kasparov's. It simply can't believe a human being would make that move—the odds are a familiar 1 in 10,000.

Just like a human, AlphaGo can be taken by surprise. Four hours and 45 minutes into the game, AlphaGo resigns. Just like us, it can lose.

"All the thinking that AlphaGo had done up to that point was sort of rendered useless," Hassabis says. "It had to restart."

The final game has begun, and I'm supposed to watch with Hassabis and his team. But just before I head to meet them, a Googler finds me in the press room. "We're so sorry," she says, "but the team has changed their mind. They don't want a reporter in the room for the final match."

After she walks away, I turn to Geordie Wood, WIRED's photographer. "You know what that means?" I say. "AlphaGo thinks it's losing."

It is. Early in the game AlphaGo makes a rookie mistake. In a crowded area on the lower half of the board, the machine places its white stone too close to Lee's line of black and loses the entire territory. AlphaGo's intuition failed it; like a human, the machine still has blind spots.

But as the game stretches into a third hour, AlphaGo claws its way back into the contest. At the three-and-a-half-hour mark, Lee's play clock runs out. Under the match rules, he now has to make each move in less than a minute or else forfeit, but a wide swath of space on the top

again, he waits until the last second to place his next stone.

Then AlphaGo's clock runs out too. Both players start moving at what looks like an impossible pace. The board fills with stones. For the first time in the series, the game looks as though it will play out to the very end—that neither side will resign before the final points are tallied. But five hours in, the gulf between Lee and AlphaGo grows too wide. Lee resigns. AlphaGo is fallible but still dominant.

In all the world, only one other person could credibly claim to know how Lee felt: Fan Hui, the three-time European champ and AlphaGo's de facto trainer. He had lost to the machine five games to nil in a closed-door match back in October, the training montage for the bigger contest in Seoul. Afterward, Fan joined DeepMind as a kind of player for hire, playing game after game against the machine—games he kept losing, one after the other.

But as Fan's losses piled up against AlphaGo, a funny thing happened. Fan came to see Go in an entirely new way. Against other humans, he started winning more—including four straight games against top players. His ranking shot up. AlphaGo was training him.

So, I ask Fan during the tournament, what should we think of Lee's fight against the machine?

The Rise of AI

Does it know who I am?

Bots that understand what they see.

They're not discerning photography critics, but the AI-powered image-recognition systems of today can do more than just see photos. They can analyze them and understand them. To hone their abilities, researchers have fed these systems kajillions of training images, a technique called deep learning. The result? Move over, Susan Sontag. Machines get photography now. And that can lead to some pretty cool abilities.
—CHELSEA LEU

Seeing for the blind
DuLight, a small earpiece-cum-camera from the Chinese search giant Baidu, hooks into the ear of a visually impaired person and pipes up about the wearer's surroundings. The gadget uses image, speech, and facial recognition to, say, tell a can of Pepsi from a can of Coke, distinguish different denominations of cash, and identify buddies.

Finding the bad stuff
In 2015 researchers at Twitter Cortex, the company's AI group, developed a system that automatically identifies NSFW images and hides them from users' streams. That way users don't have to worry about porn and beheadings filling up their #hashtags, and workers don't need to wade through scarring images to flag them manually. Cortex also uses the tech on a new job: picking out relevant tweets and posts from heaps of content.

Helping you accessorize
Pinterest's visual search tool enables its users to browse using a system trained on the company's massive trove of images. Pinners can select a particularly alluring object within a photo—say, a cast-iron pan in a table setting—and the system recommends similar pictures. Users can then use those suggestions to identify the pan's brand—and maybe even buy one.