MIT Technology Review

The predictions issue

Volume 123 Number 2
Mar/Apr 2020
USD $9.99 CAD $10.99

10 Breakthrough Technologies 2020

PLUS
How to predict what's coming in 2030 and beyond

Mega-constellations of **tiny satellites** will bring the internet to millions—and change the way we use space.

$9.99US $10.99CAN


POW!
90
YEARS SOLVING
BATTELLE
It can be done

Every year, we pick 10 recent technological breakthroughs that we predict will have a big impact in the years to come. We've been doing it for nearly two decades, and we've been pretty good at predicting big trends like data mining, natural-language processing, and microfluidics, but not so great at specific products.

Let's look back at our 2010 list: mobile phones with hologram-style 3D displays? Microbes that turn carbon dioxide from the air directly into diesel fuel? Electronic implants that dissolve in your body when their job is done? "Social TV" that lets you talk about shows with your friends online while you watch? (Yeah, we have that—it's called Twitter.)

At least in 2009 we profiled Siri—before it was even launched, mark you, let alone acquired by Apple. Shame we bought into the company's hype that it was going to be not merely a voice-activated search engine but a "do engine" that can book you a restaurant or a flight.

Then again, if we really could predict which new inventions would take off, we wouldn't tell *you* about them; we'd start a fund. Venture capitalists, who do this all day long, still get it wrong nine times out of 10. But as any decent futurist will tell you, the point of futurism isn't to guess the future; it's to challenge your assumptions about the present so the future doesn't catch you off guard.

So this year, since it's 2020 and we like round numbers as much as anyone, we decided to supplement our annual list (see page 15) with a closer look at the art and science of prediction, and to collect some other people's predictions for 2030—if only so we can have a laugh a decade hence at how wrong they were.

David Rotman (page 10) examines Moore's Law, the most reliable prediction of modern times, and asks how the predictions of its imminent demise—themselves already rather long in the tooth—will influence future progress. Rob Arthur (page 72) looks at why forecasters messed up so badly in the 2016 US presidential election and why they think they can do better in 2020. Brian Bergstein (page 62) describes the effort to create AI that understands causality so that it can make predictions more reliably. Bobbie Johnson (page 54) asks some people whose job is prediction how they think about the future and what they expect in 2030.

Meanwhile, I (page 70) pick up some more 2030 predictions at the World Economic Forum in Davos—the place where, if you believe either the conspiracy theorists or the WEF's own marketing, the future of the world is decided by politicians and billionaires. Tim Maughan (page 66) writes about design fiction, a quirky movement for imagining the future creatively, and how it got co-opted by corporations. Tate Ryan-Mosley (page 53) summarizes five big trends that will shape the next few decades, while Konstantin Kakaes (page 80) rounds up five of the best books on humanity's relationship to prediction. And Andrew Dana Hudson (page 82) provides this issue's short fiction piece, a story of one future that I fear is all too likely to come true.

We also have longer stories on some of our 10 breakthrough technologies: Erika Check Hayden on cure-for-one drugs (page 46), Ramin Skibba on satellite mega-constellations (page 30), Mike Orcutt on the future (or rather, lack thereof) of cash (page 32), and me on quantum computing (page 38).

This last topic is close to my heart; I first wrote about it more than 20 years ago, when nobody had yet built a working quantum computer. Last fall Google announced the first demonstration of "quantum supremacy," a quantum computer doing something a classical one can't feasibly pull off. Some people are still skeptical they'll ever amount to much, but I predict we will be using them to solve real problems by 2030. Check back on me then.

Gideon Lichfield is editor in chief of MIT Technology Review.

IAN ALLEN

Cover photograph by Bob O'Connor
MicroMAS-2 cubesat courtesy of MIT Lincoln Lab
Imaging: Zach Vitale

Editorial

Editor in chief
Gideon Lichfield

Executive editor
Michael Reilly

Editor at large
David Rotman

News editor
Niall Firth

Managing editor
Timothy Maher

Commissioning editors
Bobbie Johnson
Konstantin Kakaes

Senior editor, MIT News
Alice Dragoon

Senior editor, biomedicine
Antonio Regalado

Senior editor, energy
James Temple

Senior editor, ethics and policy
Angela Chen

Senior editor, cybersecurity
Patrick Howell O'Neill

Senior editor, AI
Will Douglas Heaven

Senior editor, podcasts and live journalism
Jennifer Strong

Senior reporters
Tanya Basu (humans and technology)
Karen Hao (AI)
Mike Orcutt (blockchain)

Reporters
Charlotte Jee (news)
Neel Patel (space)

Copy chief
Linda Lowenthal

Social media editor
Benji Rosen

Editorial research manager
Tate Ryan-Mosley

Administrative assistant
Andrea Siegel

Proofreader
Barbara Wallraff

Design

Chief creative officer
Eric Mongeon

Art director
Emily Luong

Marketing and events designer
Kyle Thomas Hemingway

Photo editor
Stephanie Arnett

Digital production specialist
Savash Kalay

Corporate

Chief executive officer and publisher
Elizabeth Bramson-Boudreau

Assistant to the CEO
Katie McLean

Human resources manager
James Wall

Manager of information technology
Colby Wheeler

Office manager
Linda Cardinal

Client services manager, licensing and syndication
Ted Hu

Product development

Chief digital officer
Cy Caine

Senior project manager
Allison Chase

Senior product designer
Jon Akland

Director of software engineering
Molly Frey

Director of analytics
Michelle Bellettiere

Senior software engineer
Jason Lewicki

Events

Senior vice president, events and strategic partnerships
Amy Lammers

Director of event content and experiences
Brian Bryson

Head of international and custom events
Marcy Rizzo

Event content producer
Erin Underwood

Associate director of events
Nicole Silva

Event partnership coordinator
Madeleine Frasca

Events associate
Bo Richardson

Finance

Finance director
Enejda Xheblati

General ledger manager
Olivia Male

Accountant
Letitia Trecartin

Consumer marketing

Senior vice president, marketing and consumer revenue
Doreen Adger

Director of analytics systems
Tom Russell

Director of audience development
Rosemary Kelly

Director of digital marketing
Emily Baillieul

Product marketing manager
Amanda Saeli

Assistant consumer marketing manager
Caroline da Cunha

Circulation and print production manager
Tim Borton

Advertising sales

Vice president, sales and brand partnerships
Andrew Hendler
andrew.hendler@technologyreview.com
646-520-6981

Executive director, brand partnerships
Marii Sebahar
marii@technologyreview.com
415-416-9140

Executive director, brand partnerships
Kristin Ingram
kristin.ingram@technologyreview.com
415-509-1910

Senior director, brand partnerships
Whelan Mahoney
whelan@technologyreview.com
201-417-0928

Director, brand partnerships
Debbie Hanley
debbie.hanley@technologyreview.com
214-282-2727

Director, brand partnerships
Ian Keller
ian.keller@technologyreview.com
203-858-3396

Business development sales manager
Ken Collina
ken.collina@technologyreview.com
617-475-8004

Advertising services
webcreative@technologyreview.com
617-475-8004

Media kit
www.technologyreview.com/media

MIT Technology Review Insights

Vice president of international business development, head of MIT Technology Review Insights
Nicola Crepaldi

Director of custom content, US
Laurel Ruma

Content manager
Jason Sparapani

Senior project manager
Martha Leibs

Director of custom content, international
Claire Beatty

Director of business development, Asia
Marcus Ulvne

Customer service and subscription inquiries

National
800-877-5230

International
903-636-1115

Email
customer_service@mittechnologyreview.info

Web
www.technologyreview.com/customerservice

MIT Records (alums only)
617-253-8270

Reprints
techreview@wrightsmedia.com
877-652-5295

Licensing and permissions
licensing@technologyreview.com

MIT Technology Review

One Main Street
13th Floor
Cambridge, MA 02142
617-475-8000

The mission of MIT Technology Review is to make technology a greater force for good by bringing about better-informed, more conscious technology decisions through authoritative, influential, and trustworthy journalism.

Technology Review, Inc., is an independent nonprofit 501(c)(3) corporation wholly owned by MIT; the views expressed in our publications and at our events are not always shared by the Institute.


March 23–25, 2020
San Francisco, CA
MIT Technology Review
EmTech Digital
Are you AI ready?
EmTech Digtial offers executive decision makers ideas and inspiration addressing essential technical, ethical and strategic AI deployment issues facing organizations in 2020.
Join us as we examine AI topics such as trust, research, algorithms, convergence, workforce, and strategy.
83%
EmTech Digital

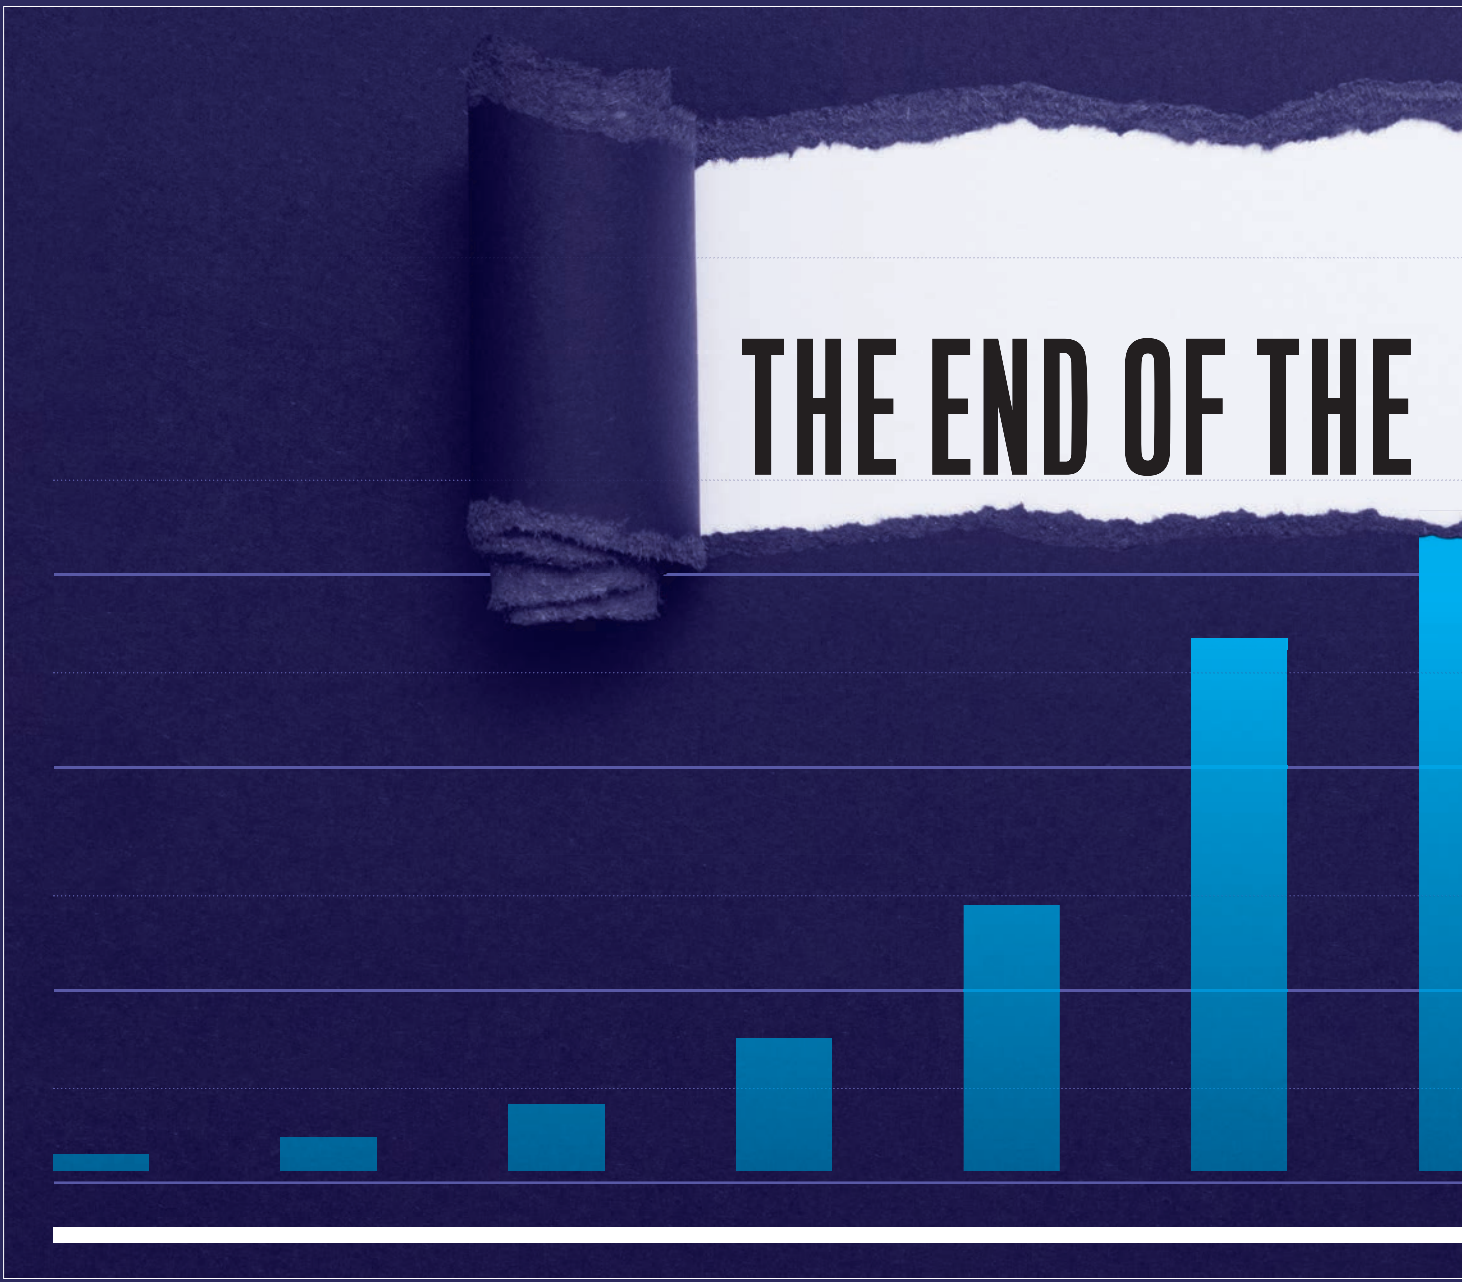

THE END OF THE

MOORE'S LAW FUELED PROSPERITY FOR 50 YEARS, BUT IT'S ENDING. WE HAVE NO IDEA WHAT COMES NEXT.

BY DAVID ROTMAN

Gordon Moore's 1965 forecast that the number of components on an integrated circuit would double every year until it reached an astonishing 65,000 by 1975 is the greatest technological prediction of the last half-century. When it proved correct in 1975, he revised what has become known as Moore's Law to a doubling of transistors on a chip every two years.

Since then, his prediction has defined the trajectory of technology and, in many ways, of progress itself.

Moore's argument was an economic one. Integrated circuits, with multiple transistors and other electronic devices interconnected with aluminum metal lines on a tiny square of silicon wafer, had been invented a few years earlier by Robert Noyce at Fairchild Semiconductor. Moore, the company's R&D director,

GREATEST PREDICTION ON EARTH

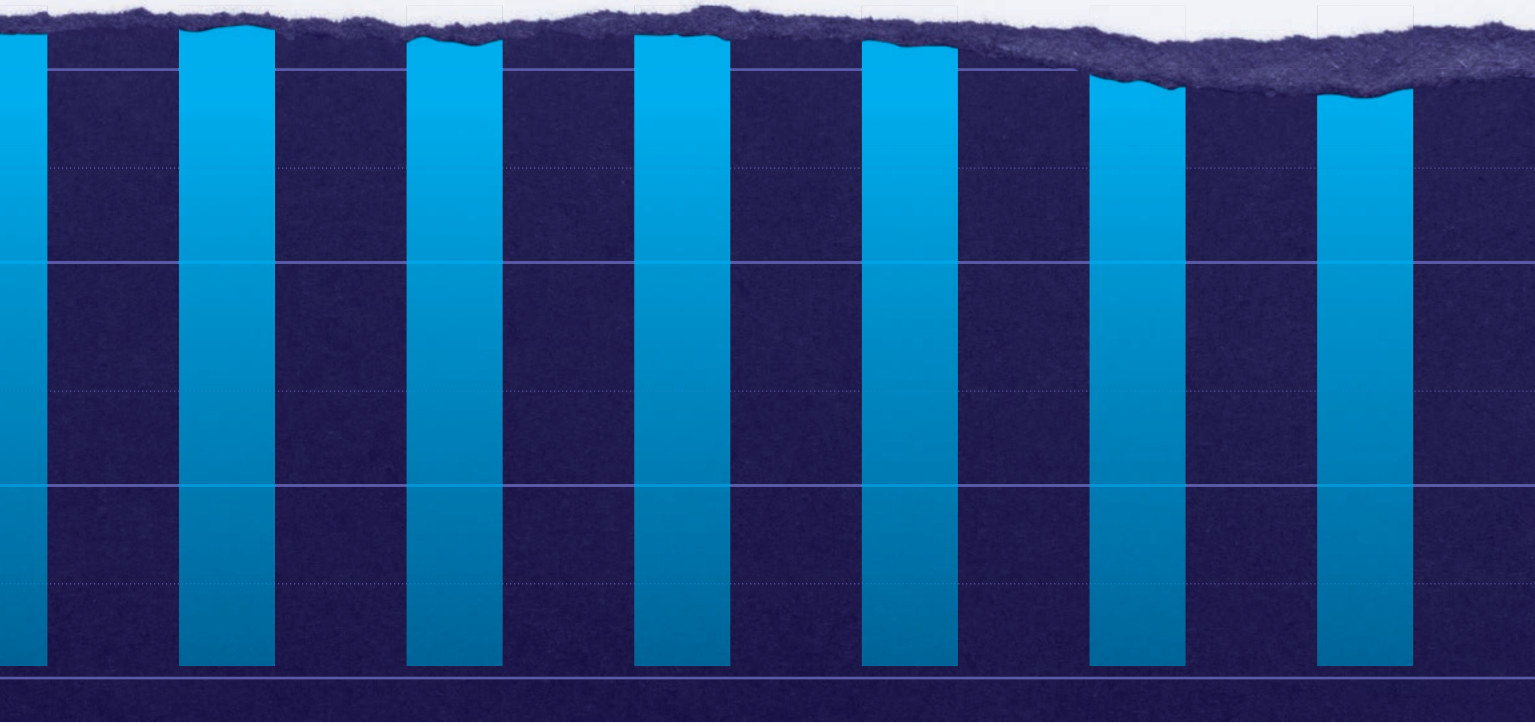

realized, as he wrote in 1965, that with these new integrated circuits, “the cost per component is nearly inversely proportional to the number of components.” It was a beautiful bargain—in theory, the more transistors you added, the cheaper each one got. Moore also saw that there was plenty of room for engineering advances to increase the number of transistors you could affordably and reliably put on a chip.

Soon these cheaper, more powerful chips would become what economists like to call a general purpose technology—one so fundamental that it spawns all sorts of other innovations and advances in multiple industries. A few years ago, leading economists credited the information technology made possible by integrated circuits with a third of US productivity growth since 1974. Almost every technology we care about, from smartphones to cheap laptops to GPS, is a direct reflection of Moore’s prediction. It has also fueled today’s breakthroughs in artificial intelligence and genetic medicine, by giving machine-learning techniques the ability to chew through massive amounts of data to find answers.

But how did a simple prediction, based on extrapolating from a graph of the number of transistors by year—a graph that at the time had only a few data points—come to define a half-century of progress?

In part, at least, because the semiconductor industry decided it would.

Moore wrote that "cramming more components onto integrated circuits," the title of his 1965 article, would "lead to such wonders as home computers—or at least terminals connected to a central computer—automatic controls for automobiles, and personal portable communications equipment." In other words, stick to his road map of squeezing ever more transistors onto chips and it would lead you to the promised land. And for the following decades, a booming industry, the government, and armies of academic and industrial researchers poured money and time into upholding Moore's Law, creating a self-fulfilling prophecy that kept progress on track with uncanny accuracy. Though the pace of progress has slipped in recent years, the most advanced chips today have nearly 50 billion transistors.

Every year since 2001, MIT Technology Review has chosen the 10 most important breakthrough technologies of the year. It's a list of technologies that, almost without exception, are possible only because of the computation advances described by Moore's Law.

For some of the items on this year's list the connection is obvious: consumer devices, including watches and phones, infused with AI; climate-change attribution made possible by improved computer modeling and data gathered from worldwide atmospheric monitoring systems; and cheap, pint-size satellites. Others on the list, including quantum supremacy, molecules discovered using AI, and even anti-aging treatments and hyper-personalized drugs, are due largely to the computational power available to researchers.

But what happens when Moore's Law inevitably ends? Or what if, as some suspect, it has already died, and we are already running on the fumes of the greatest technology engine of our time?

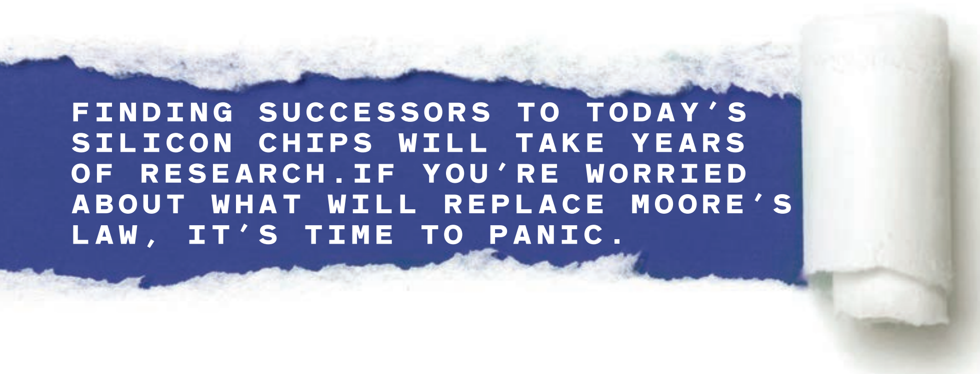

RIP

"It's over. This year that became really clear," says Charles Leiserson, a computer scientist at MIT and a pioneer of parallel computing, in which multiple calculations are performed simultaneously. The newest Intel fabrication plant, meant to build chips with minimum feature sizes of 10 nanometers, was much delayed, delivering chips in 2019, five years after the previous generation of chips with 14-nanometer features. Moore's Law, Leiserson says, was always about the rate of progress, and "we're no longer on that rate." Numerous other prominent computer scientists have also declared Moore's Law dead in recent years. In early 2019, the CEO of the large chipmaker Nvidia agreed.

In truth, it's been more a gradual decline than a sudden death. Over the decades, some, including Moore himself at times, fretted that they could see the end in sight, as it got harder to make smaller and smaller transistors. In 1999, an Intel researcher worried that the industry's goal of making transistors smaller than 100 nanometers by 2005 faced fundamental physical problems with "no known solutions," like the quantum effects of electrons wandering where they shouldn't be.

For years the chip industry managed to evade these physical roadblocks. New transistor designs were introduced to better corral the electrons. New lithography methods using extreme ultraviolet radiation were invented when the wavelengths of visible light were too thick to precisely carve out silicon features of only a few tens of nanometers. But progress grew ever more expensive. Economists at Stanford and MIT have calculated that the research effort going into upholding Moore's Law has risen by a factor of 18 since 1971.

Likewise, the fabs that make the most advanced chips are becoming prohibitively pricey. The cost of a fab is rising at around 13% a year, and is expected to reach $16 billion or more by 2022. Not coincidentally, the number of companies with plans to make the next generation of chips has now shrunk to only three, down from eight in 2010 and 25 in 2002.

Nonetheless, Intel—one of those three chipmakers—isn't expecting a funeral for Moore's Law anytime soon. Jim Keller, who took over as Intel's head of silicon engineering in 2018, is the man with the job of keeping it alive. He leads a team of some 8,000 hardware engineers and chip designers at Intel. When he joined the company, he says, many were anticipating the end of Moore's Law. If they were right, he recalls thinking, "that's a drag" and maybe he had made "a really bad career move."

But Keller found ample technical opportunities for advances. He points out that there are probably more than a hundred variables involved in keeping Moore's Law going, each of which provides different benefits and faces its own limits. It means there are many ways to keep doubling the number of devices on a chip—innovations such as 3D architectures and new transistor designs.

GETTY IMAGES

These days Keller sounds optimistic. He says he has been hearing about the end of Moore's Law for his entire career. After a while, he "decided not to worry about it." He says Intel is on pace for the next 10 years, and he will happily do the math for you: 65 billion (number of transistors) times 32 (if chip density doubles every two years) is 2 trillion transistors. "That's a 30 times improvement in performance," he says, adding that if software developers are clever, we could get chips that are a hundred times faster in 10 years.

Still, even if Intel and the other remaining chipmakers can squeeze out a few more generations of even more advanced microchips, the days when you could reliably count on faster, cheaper chips every couple of years are clearly over. That doesn't, however, mean the end of computational progress.

Time to panic

Neil Thompson is an economist, but his office is at CSAIL, MIT's sprawling AI and computer center, surrounded by roboticists and computer scientists, including his collaborator Leiserson. In a new paper, the two document ample room for improving computational performance through better software, algorithms, and specialized chip architecture.

One opportunity is in slimming down so-called software bloat to wring the most out of existing chips. When chips could always be counted on to get faster and more powerful, programmers didn't need to worry much about writing more efficient code. And they often failed to take full advantage of changes in hardware architecture, such as the multiple cores, or processors, seen in chips used today.

Thompson and his colleagues showed that they could get a computationally intensive calculation to run some 47 times faster just by switching from Python, a popular general-purpose programming language, to the more efficient C. That's because C, while it requires more work from the programmer, greatly reduces the required number of operations, making a program run much faster. Further tailoring the code to take full advantage of a chip with 18 processing cores sped things up even more. In just 0.41 seconds, the researchers got a result that took seven hours with Python code.

That sounds like good news for continuing progress, but Thompson worries it also signals the decline of computers as a general purpose technology. Rather than "lifting all boats," as Moore's Law has, by offering ever faster and cheaper chips that were universally available, advances in software and specialized architecture will now start to selectively target specific problems and business opportunities, favoring those with sufficient money and resources.

Indeed, the move to chips designed for specific applications, particularly in AI, is well under way. Deep learning and other AI applications increasingly rely on graphics processing units (GPUs) adapted from gaming, which can handle parallel operations, while companies like Google, Microsoft, and Baidu are designing AI chips for their own particular needs. AI, particularly deep learning, has a huge appetite for computer power, and specialized chips can greatly speed up its performance, says Thompson.

But the trade-off is that specialized chips are less versatile than traditional CPUs. Thompson is concerned that chips for more general computing are becoming a backwater, slowing "the overall pace of computer improvement," as he writes in an upcoming paper, "The Decline of Computers as a General Purpose Technology."

At some point, says Erica Fuchs, a professor of engineering and public policy at Carnegie Mellon, those developing AI and other applications will miss the decreases in cost and increases in performance delivered by Moore's Law. "Maybe in 10 years or 30 years—no one really knows when—you're going to need a device with that additional computation power," she says.

The problem, says Fuchs, is that the successors to today's general purpose chips are unknown and will take years of basic research and development to create. If you're worried about what will replace Moore's Law, she suggests, "the moment to panic is now." There are, she says, "really smart people in AI who aren't aware of the hardware constraints facing long-term advances in computing." What's more, she says, because application-specific chips are proving hugely profitable, there are few incentives to invest in new logic devices and ways of doing computing.

Wanted: A Marshall Plan for chips

In 2018, Fuchs and her CMU colleagues Hassan Khan and David Hounshell wrote a paper tracing the history of Moore's Law and identifying the changes behind today's lack of the industry and government collaboration that fostered so much progress in earlier decades. They argued that "the splintering of the technology trajectories and the short-term private profitability of many of these new splinters" means we need to greatly boost public investment in finding the next great computer technologies.

If economists are right, and much of the growth in the 1990s and early 2000s was a result of microchips—and if, as some suggest, the sluggish productivity growth that began in the mid-2000s reflects the slowdown in computational progress—then, says Thompson, "it follows you should invest enormous amounts of money to find the successor technology. We're not doing it. And it's a public policy failure."

There's no guarantee that such investments will pay off. Quantum computing, carbon nanotube transistors, even spintronics, are enticing possibilities—but none are obvious replacements for the promise that Gordon Moore first saw in a simple integrated circuit. We need the research investments now to find out, though. Because one prediction is pretty much certain to come true: we're always going to want more computing power.

David Rotman is editor at large of MIT Technology Review.

10 BREAKTHROUGH TECHNOLOGIES 2020

Here is our annual list of technological advances that we believe will make a real difference in solving important problems. How do we pick? We avoid the one-off tricks, the overhyped new gadgets. Instead we look for those breakthroughs that will truly change how we live and work.

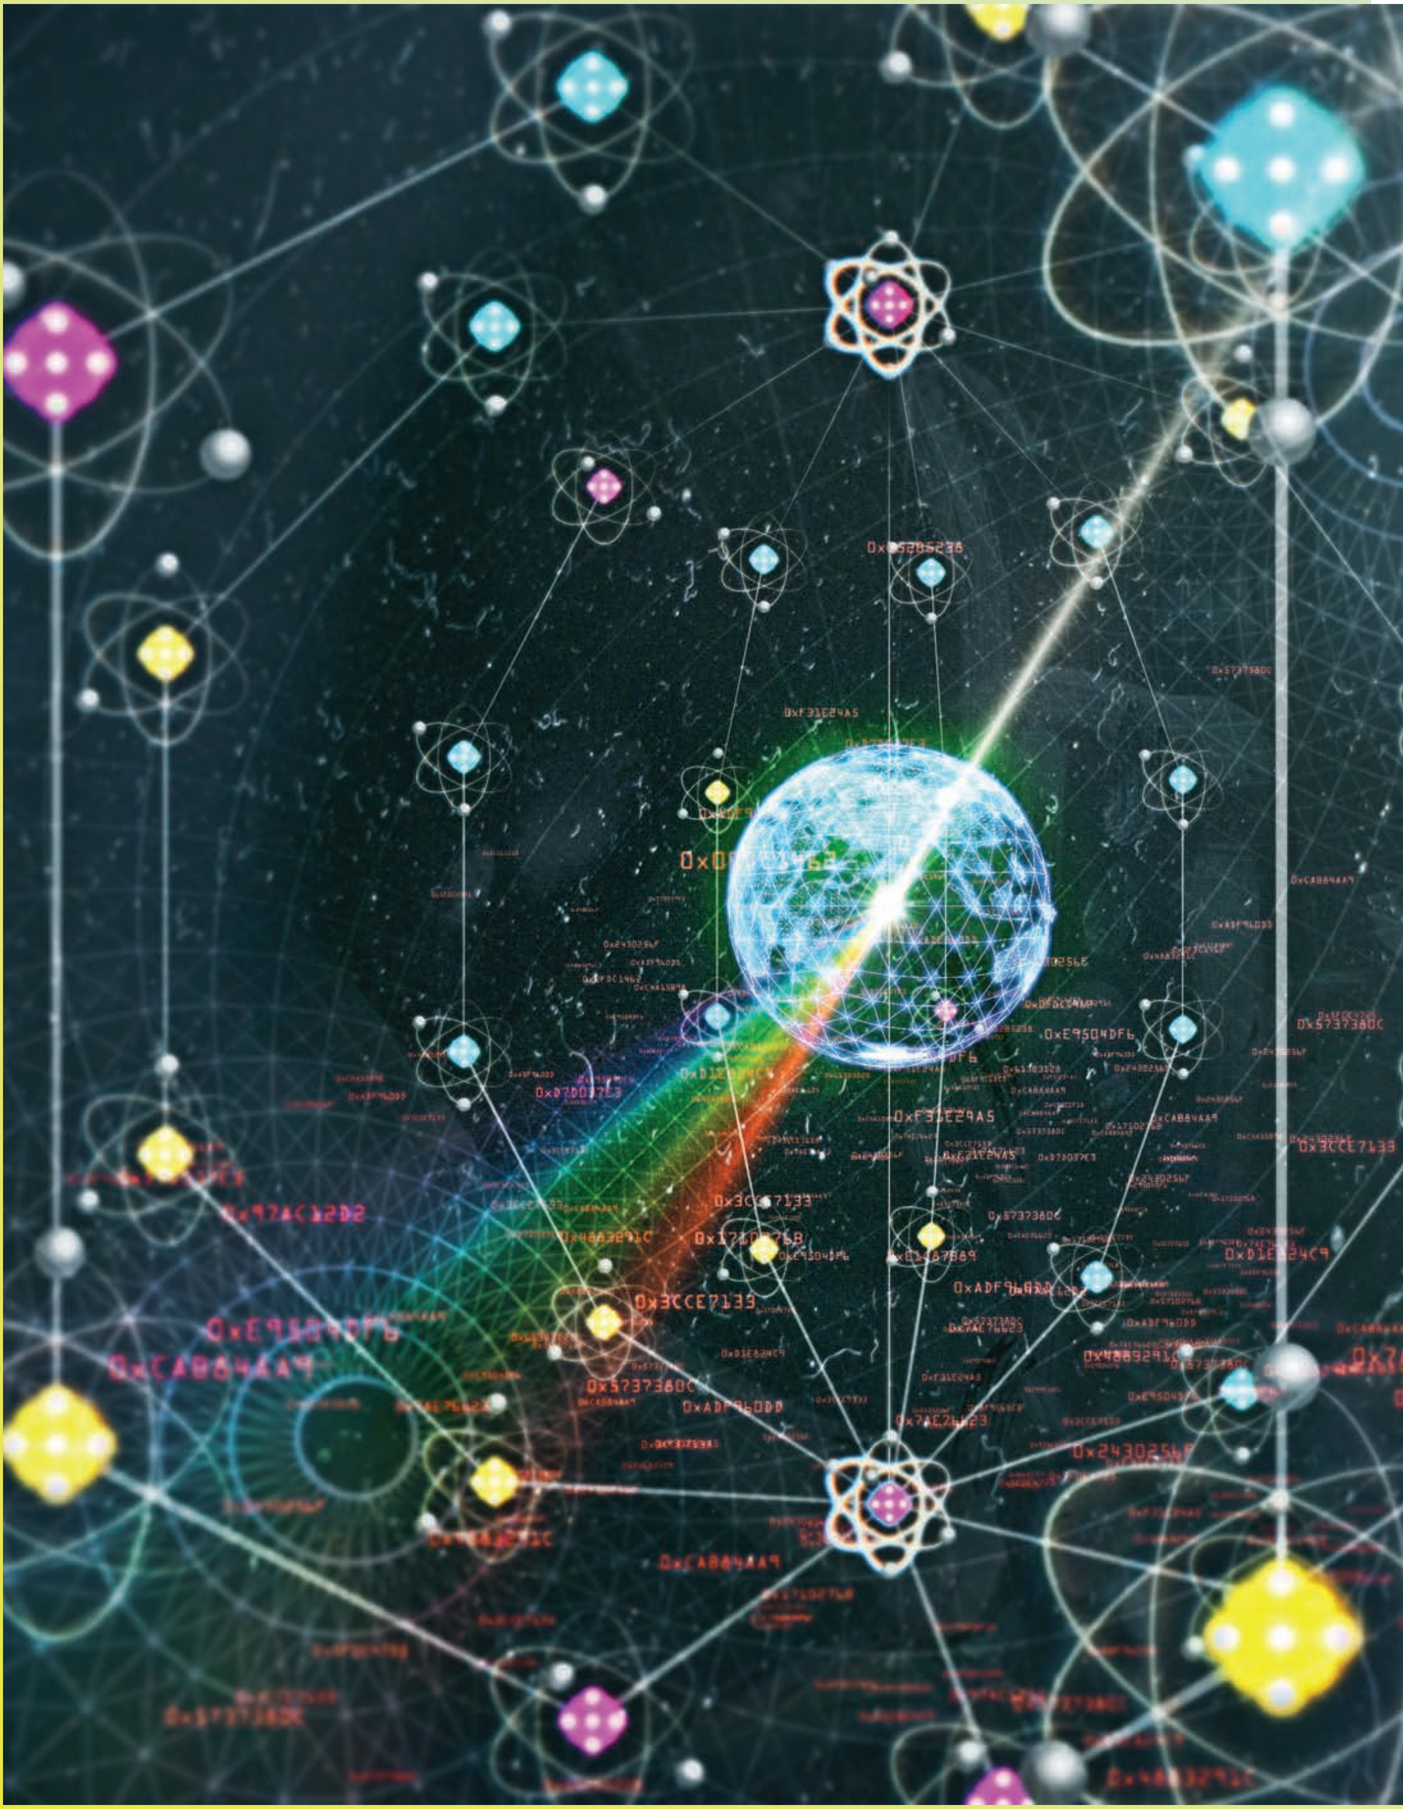

YOSHI SODEOKA

Later this year, Dutch researchers will complete a quantum internet between Delft and the Hague.

UNHACKABLE INTERNET

WHY IT MATTERS

The internet is increasingly vulnerable to hacking; a quantum one would be unhackable.

KEY PLAYERS

Delft University of Technology

Quantum Internet Alliance

University of Science and Technology of China

AVAILABILITY

5 years

An internet based on quantum physics will soon enable inherently secure communication. A team led by Stephanie Wehner, at Delft University of Technology, is building a network connecting four cities in the Netherlands entirely by means of quantum technology. Messages sent over this network will be unhackable.

In the last few years, scientists have learned to transmit pairs of photons across fiber-optic cables in a way that absolutely protects the information encoded in them. A team in China used a form of the technology to construct a 2,000-kilometer network backbone between Beijing and Shanghai—but that project relies partly on classical components that periodically break the quantum link before establishing a new one, introducing the risk of hacking.

The Delft network, in contrast, will be the first to transmit information between cities using quantum techniques from end to end.

The technology relies on a quantum behavior of atomic particles called entanglement. Entangled photons can't be covertly read without disrupting their content.

But entangled particles are difficult to create, and harder still to transmit over long distances. Wehner's team has demonstrated it can send them more than 1.5 kilometers (0.93 miles), and they are confident they can set up a quantum link between Delft and the Hague by around the end of this year. Ensuring an unbroken connection over greater distances will require quantum repeaters that extend the network.

Such repeaters are currently in design at Delft and elsewhere. The first should be completed in the next five to six years, says Wehner, with a global quantum network following by the end of the decade. —*Russ Juskalian*

HYPER-PERSONALIZED MEDICINE

Novel drugs are being designed to treat

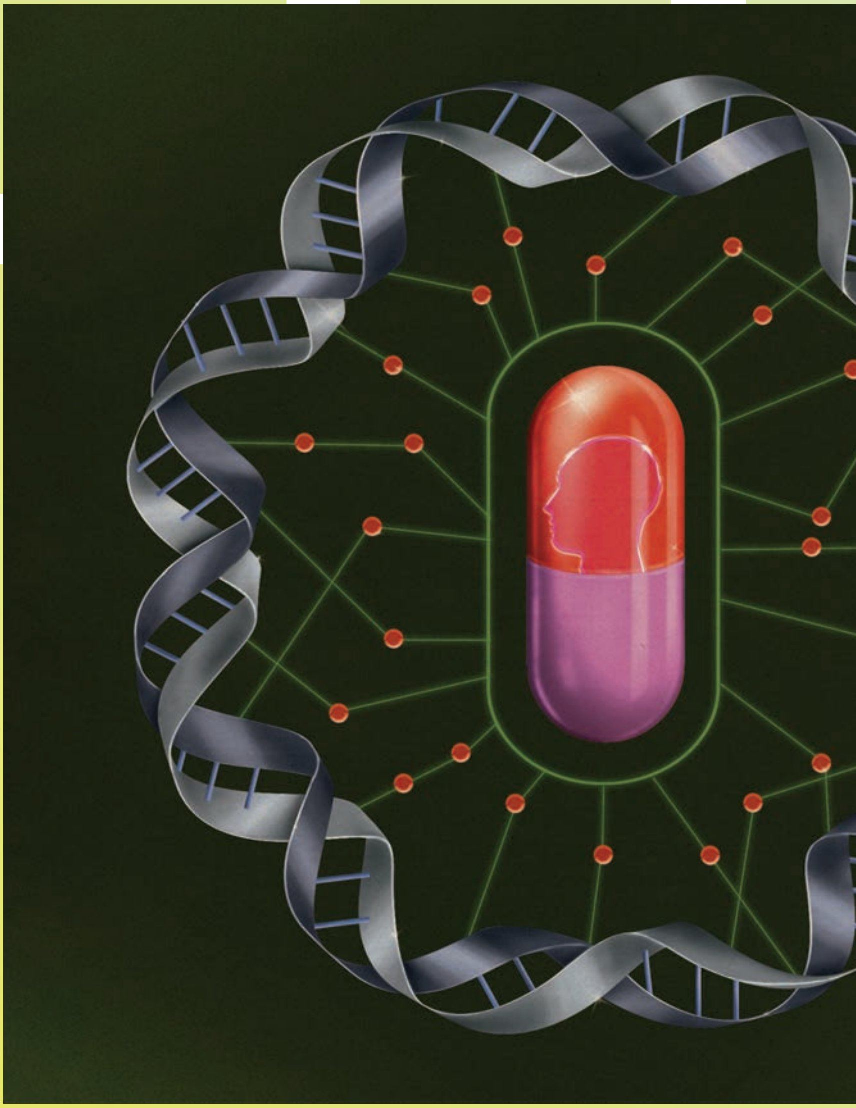

Here's a definition of a hopeless case: a child with a fatal disease so exceedingly rare that not only is there no treatment, there's not even anyone in a lab coat studying it. "Too rare to care," goes the saying.

That's about to change, thanks to new classes of drugs that can be tailored to a person's genes. If an extremely rare disease is caused by a specific DNA mistake—as several thousand are—there's now at least a fighting chance for a genetic fix.

One such case is that of Mila Makovec, a little girl suffering from a devastating illness caused by a unique genetic mutation, who got a drug manufactured just for her. Her case made the New England Journal of Medicine in October, after doctors moved from a readout of her genetic error to a treatment in just a year. They called the drug milasen, after her.

The treatment hasn't cured Mila. But it seems to have stabilized her condition: it has reduced her seizures, and she has begun to stand and walk with assistance.

Mila's treatment was possible because creating a gene medicine has never been faster or had a better chance of working. The new medicines might take the form of gene replacement, gene editing, or antisense (the type Mila received), a sort of molecular eraser, which erases or fixes erroneous genetic messages. What the treatments have in common is that they can be programmed, in digital fashion and with digital speed, to correct or compensate for inherited diseases, letter for DNA letter.

How many stories like Mila's are there? So far, just a handful.

But more are on the way. Where researchers would have once seen obstacles and said "I'm sorry," they now see solutions in DNA and think maybe they can help.

The real challenge for "n-of-1" treatments (a reference to the number of people who get the drug) is that they defy just about every accepted notion of how pharmaceuticals should be developed, tested, and sold. Who will pay for these drugs when they help one person, but still take large teams to design and manufacture?

—Antonio Regalado

unique genetic mutations.

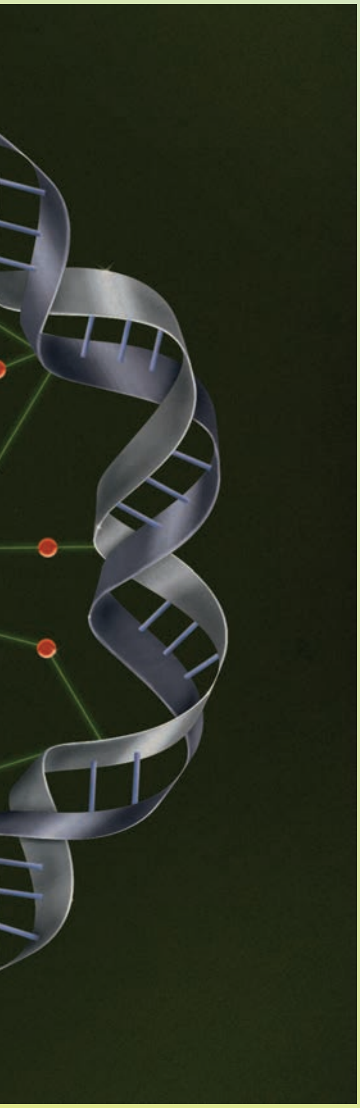

JULIA DUFOSSÉ

WHY IT MATTERS

Genetic medicine tailored to a single patient means hope for people whose ailments were previously uncurable.

KEY PLAYERS

A-T Children's Project

Boston Children's Hospital

Ionis Pharmaceuticals

US Food & Drug Administration

AVAILABILITY

Now

DIGITAL MONEY

The rise of digital currency has massive ramifications for financial privacy.

WHY IT MATTERS

As the use of physical cash declines, so does the freedom to transact without an intermediary. Meanwhile, digital currency technology could be used to splinter the global financial system.

KEY PLAYERS

People's Bank of China

Facebook

AVAILABILITY

This year

Last June Facebook unveiled a "global digital currency" called Libra. The idea triggered a backlash and Libra may never launch, at least not in the way it was originally envisioned. But it's still made a difference: just days after Facebook's announcement, an official from the People's Bank of China implied that it would speed the development of its own digital currency in response. Now China is poised to become the first major economy to issue a digital version of its money, which it intends as a replacement for physical cash.

China's leaders apparently see Libra, meant to be backed by a reserve that will be mostly US dollars, as a threat: it could reinforce America's disproportionate power over the global financial system, which stems from the dollar's role as the world's de facto reserve currency. Some suspect China intends to promote its digital renminbi internationally.

Now Facebook's Libra pitch has become geopolitical. In October, CEO Mark Zuckerberg promised Congress that Libra "will extend America's financial leadership as well as our democratic values and oversight around the world." The digital money wars have begun. —*Mike Orcutt*

Drugs that try to treat ailments by targeting a natural aging process in the body have shown promise.

ANTI-AGING DRUGS

The first wave of a new class of anti-aging drugs have begun human testing. These drugs won't let you live longer (yet) but aim to treat specific ailments by slowing or reversing a fundamental process of aging.

The drugs are called senolytics—they work by removing certain cells that accumulate as we age. Known as "senescent" cells, they can create low-level inflammation that suppresses normal mechanisms of cellular repair and creates a toxic environment for neighboring cells.

In June, San Francisco–based Unity Biotechnology reported initial results in patients with mild to severe osteoarthritis of the knee. Results from a larger clinical trial are expected in the second half of 2020. The company is also developing similar drugs to treat age-related diseases of the eyes and lungs, among other conditions.

Senolytics are now in human tests, along with a number of other promising approaches targeting the biological processes that lie at the root of aging and various diseases.

A company called Alkahest injects patients with components found in young people's blood and says it hopes to halt cognitive and functional decline in patients suffering from mild to moderate Alzheimer's disease. The company also has drugs for Parkinson's and dementia in human testing.

And in December, researchers at Drexel University College of Medicine even tried to see if a cream including the immune-suppressing drug rapamycin could slow aging in human skin.

The tests reflect researchers' expanding efforts to learn if the many diseases associated with getting older—such as heart diseases, arthritis, cancer, and dementia—can be hacked to delay their onset.

—*Adam Piore*

WHY IT MATTERS

A number of different diseases, including cancer, heart disease, and dementia, could potentially be treated by slowing aging.

KEY PLAYERS

Unity Biotechnology

Alkahest

Mayo Clinic

Oisín Biotechnologies

AVAILABILITY

Less than 5 years

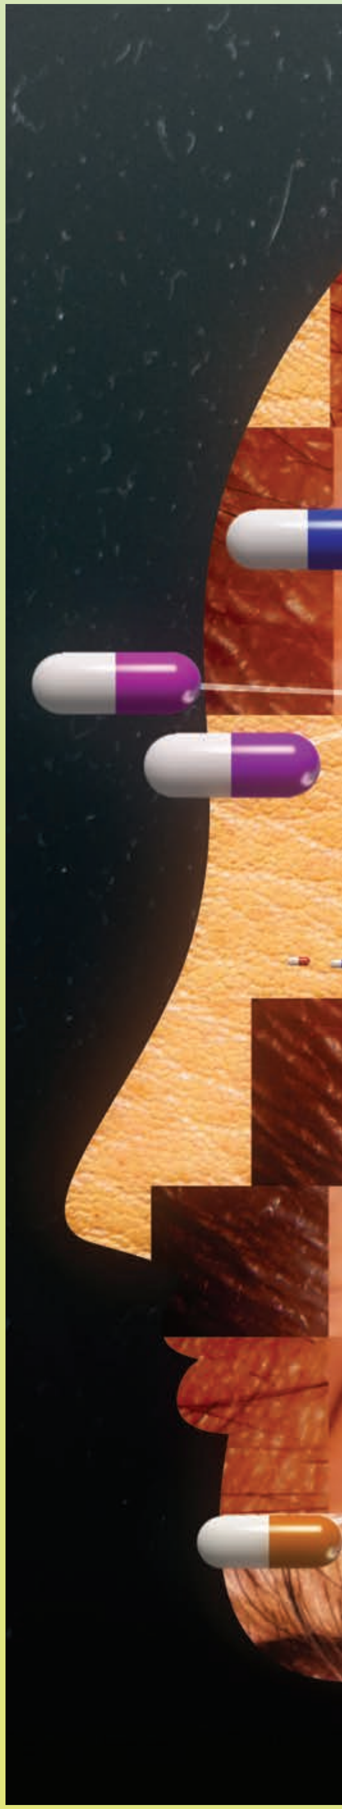

YOSHI SODEOKA

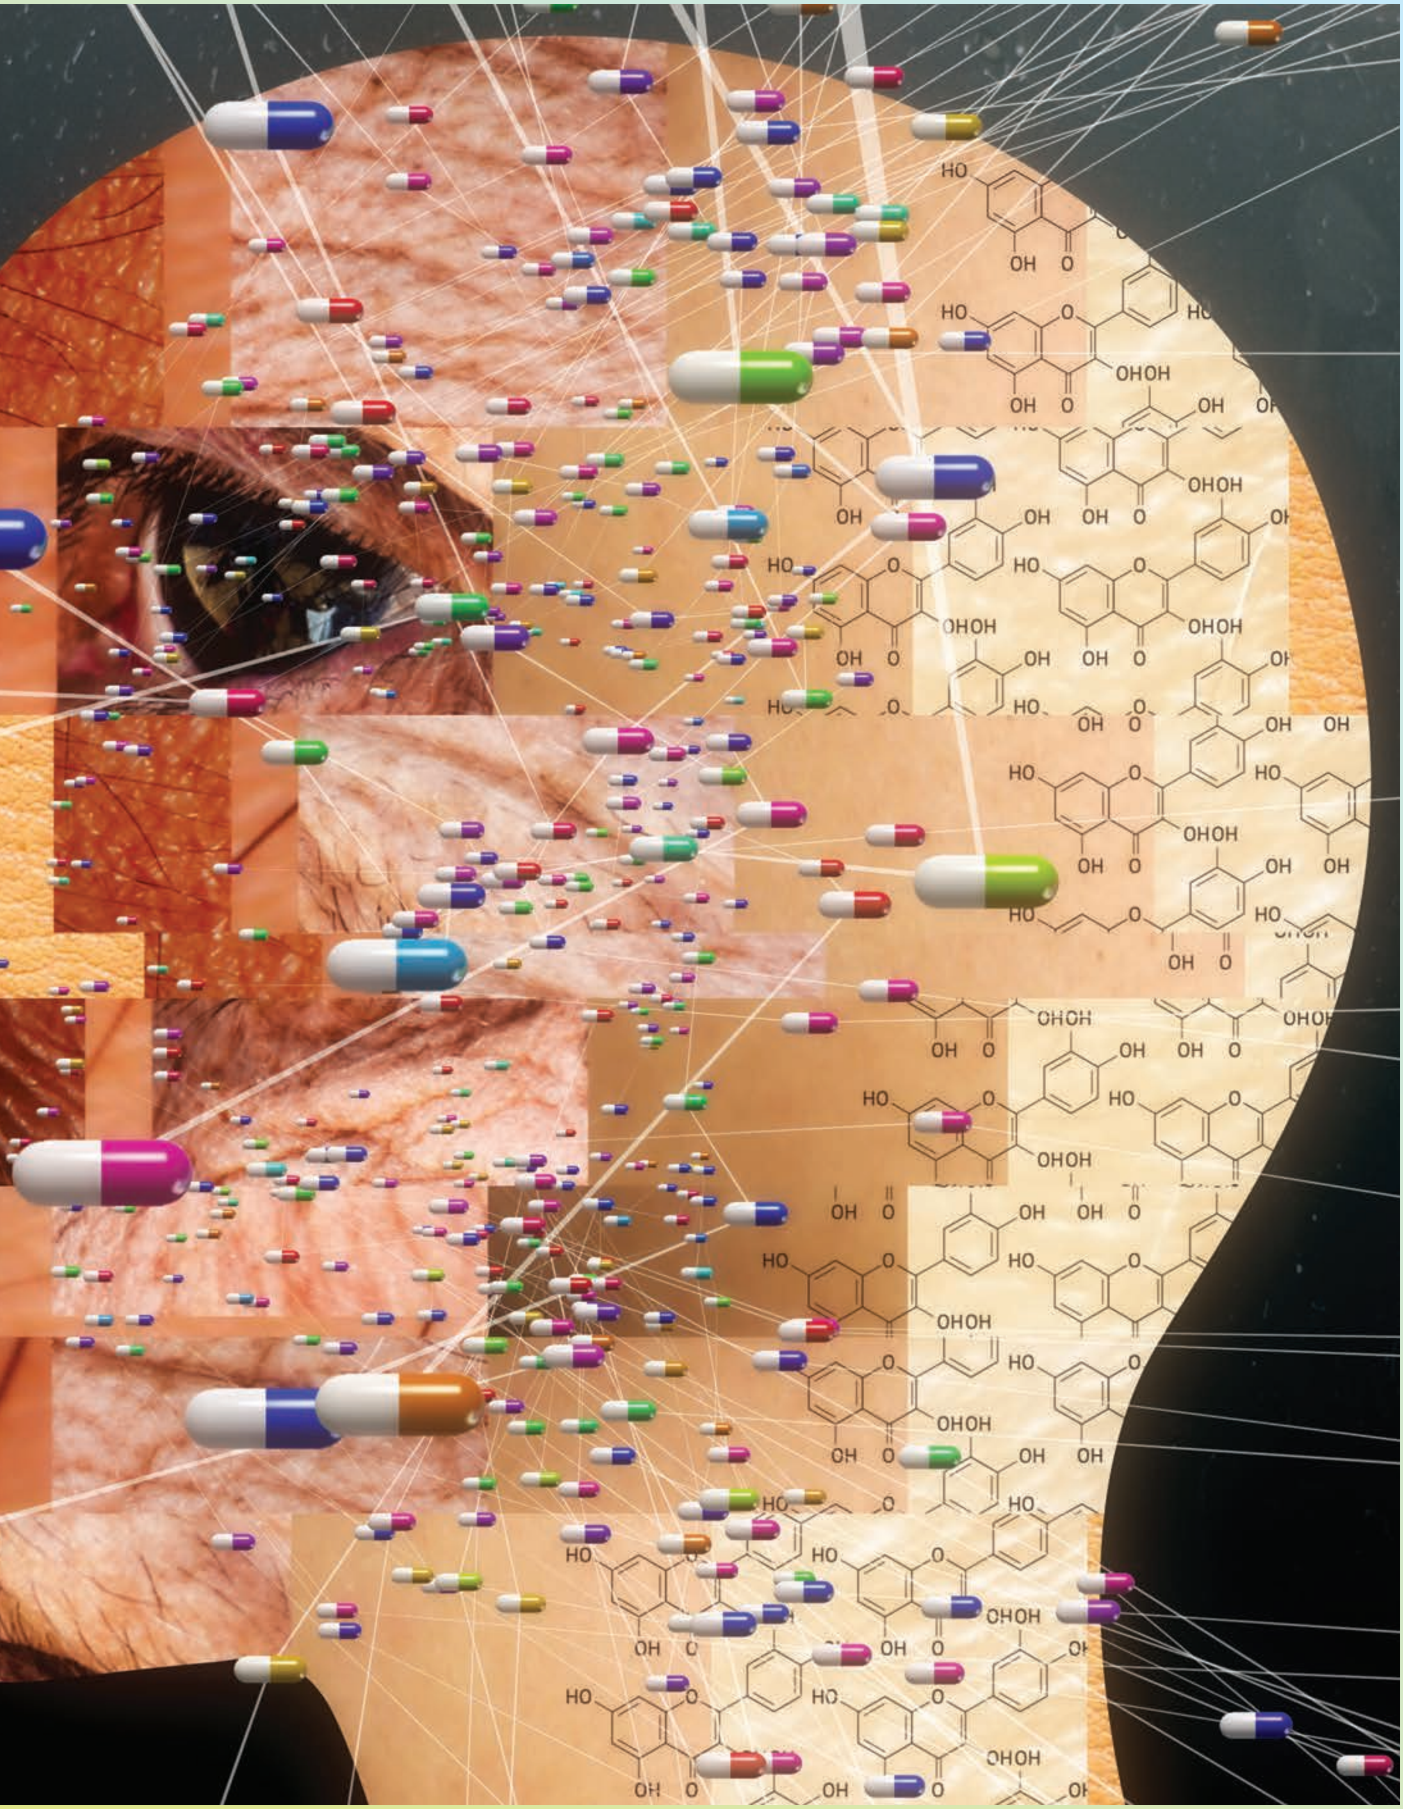

AI-DISCOVERED MOLECULES

Scientists have used AI to discover promising drug-like compounds.

WHY IT MATTERS

Commercializing a new drug costs around $2.5 billion on average. One reason is the difficulty of finding promising molecules.

KEY PLAYERS

Insilico Medicine

Kebotix

Atomwise

University of Toronto

BenevolentAI

AVAILABILITY

3-5 years

The universe of molecules that could be turned into potentially life-saving drugs is mind-boggling in size: researchers estimate the number at around 10^{60}. That's more than all the atoms in the solar system, offering virtually unlimited chemical possibilities—if only chemists could find the worthwhile ones.

Now machine-learning tools can explore large databases of existing molecules and their properties, using the information to generate new possibilities. This could make it faster and cheaper to discover new drug candidates.

In September, a team of researchers at Hong Kong–based Insilico Medicine and the University of Toronto took a convincing step toward showing that the strategy works by synthesizing several drug candidates found by AI algorithms.

Using techniques like deep learning and generative models similar to the ones that allowed a computer to beat the world champion at the ancient game of Go, the researchers identified some 30,000 novel molecules with desirable properties. They selected six to synthesize and test. One was particularly active and proved promising in animal tests.

Chemists in drug discovery often dream up new molecules—an art honed by years of experience and, among the best drug hunters, by a keen intuition. Now these scientists have a new tool to expand their imaginations.

—*David Rotman*

10

SATELLITE

More than 3.5 billion people in the world still lack internet access. Companies like SpaceX and OneWeb think they can connect every inch of the planet by launching mega-constellations of thousands of satellites that can beam a broadband connection to internet terminals. As long as these terminals have a clear view of the sky, they can deliver internet to any nearby devices. SpaceX alone wants to send more than 4.5 times more satellites into orbit this decade than humans have ever launched since *Sputnik*.

These mega-constellations are feasible because we have learned how to build smaller satellites and launch them more cheaply. During the space shuttle era, launching a satellite into space cost roughly $24,800 per pound. A small communications satellite that weighed four tons cost nearly $200 million to fly up.

Today a SpaceX Starlink satellite weighs about 500 pounds (227 kilograms). Reusable architecture and cheaper manufacturing mean we can strap dozens of them onto rockets to greatly lower the cost; a SpaceX Falcon 9 launch today costs about $1,240 per pound.

The first 120 Starlink satellites went up last year, and the company planned to launch batches of 60 every two weeks starting in January 2020. OneWeb will launch over 30 satellites later this year. We could soon see thousands of satellites working in tandem to supply internet access for even the poorest and most remote populations on the planet.

But that's only if things work out. Some researchers are livid because they fear these objects will disrupt astronomy research. Worse is the prospect of a collision that could cascade into a catastrophe of millions of pieces of space debris, making satellite services and future space exploration next to impossible. Starlink's near-miss with an ESA weather satellite in September was a jolting reminder that the world is woefully unprepared to manage this much orbital traffic. What happens with these mega-constellations this decade will define the future of orbital space. —*Neel V. Patel*

JULIA DUFOSSÉ

We can now affordably build, launch, and operate tens of thousands of satellites in orbit at once.

MEGA-CONSTELLATIONS

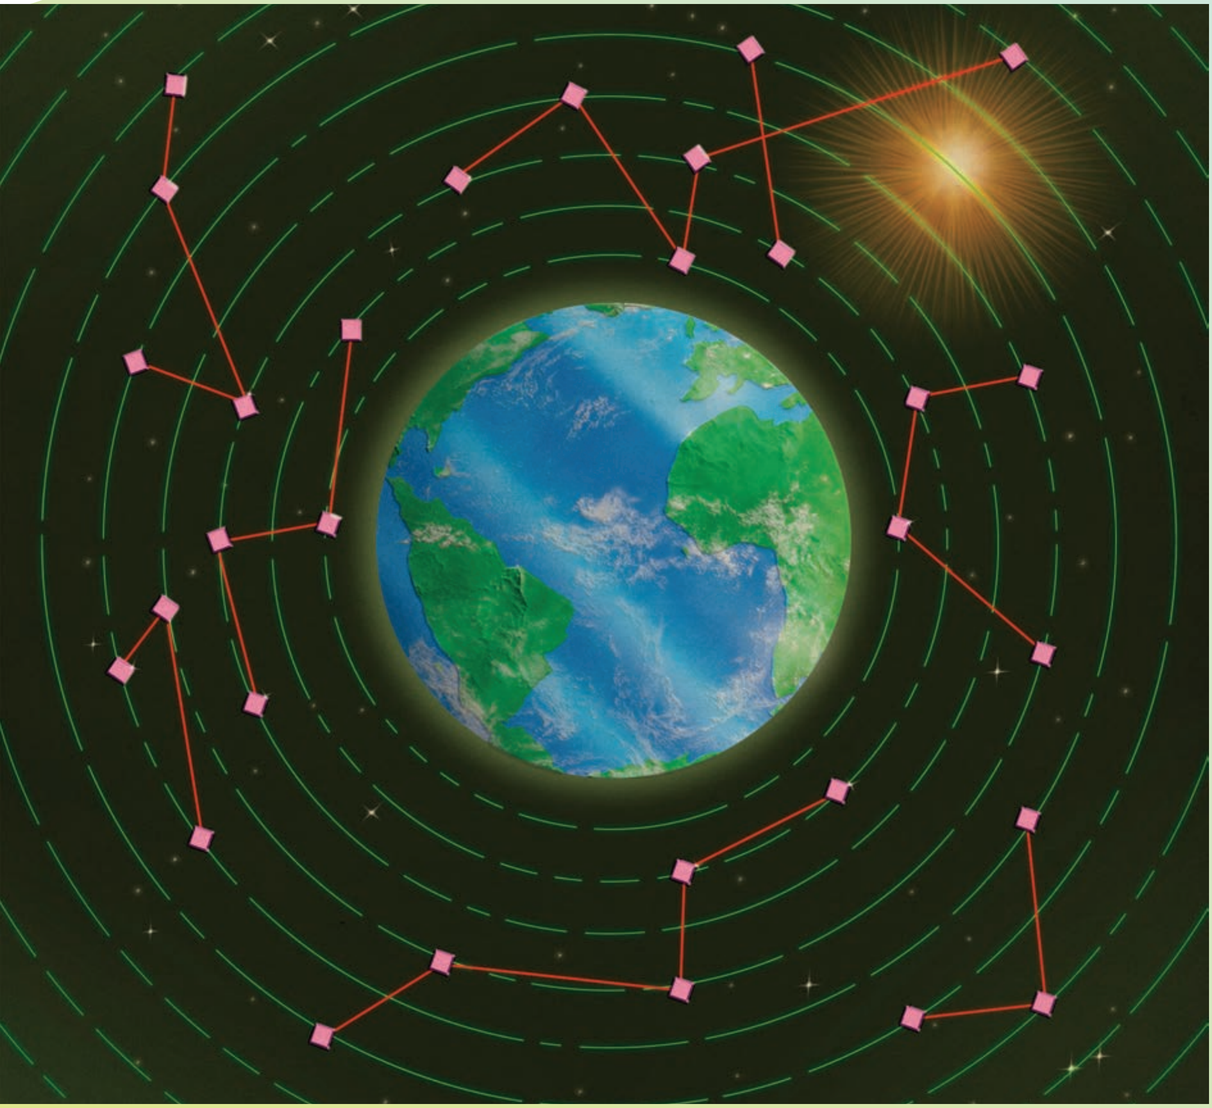

WHY IT MATTERS

These systems can blanket the globe with high-speed internet—or turn Earth's orbit into a junk-ridden minefield.

KEY PLAYERS

SpaceX

OneWeb

Amazon

Telesat

AVAILABILITY

Now

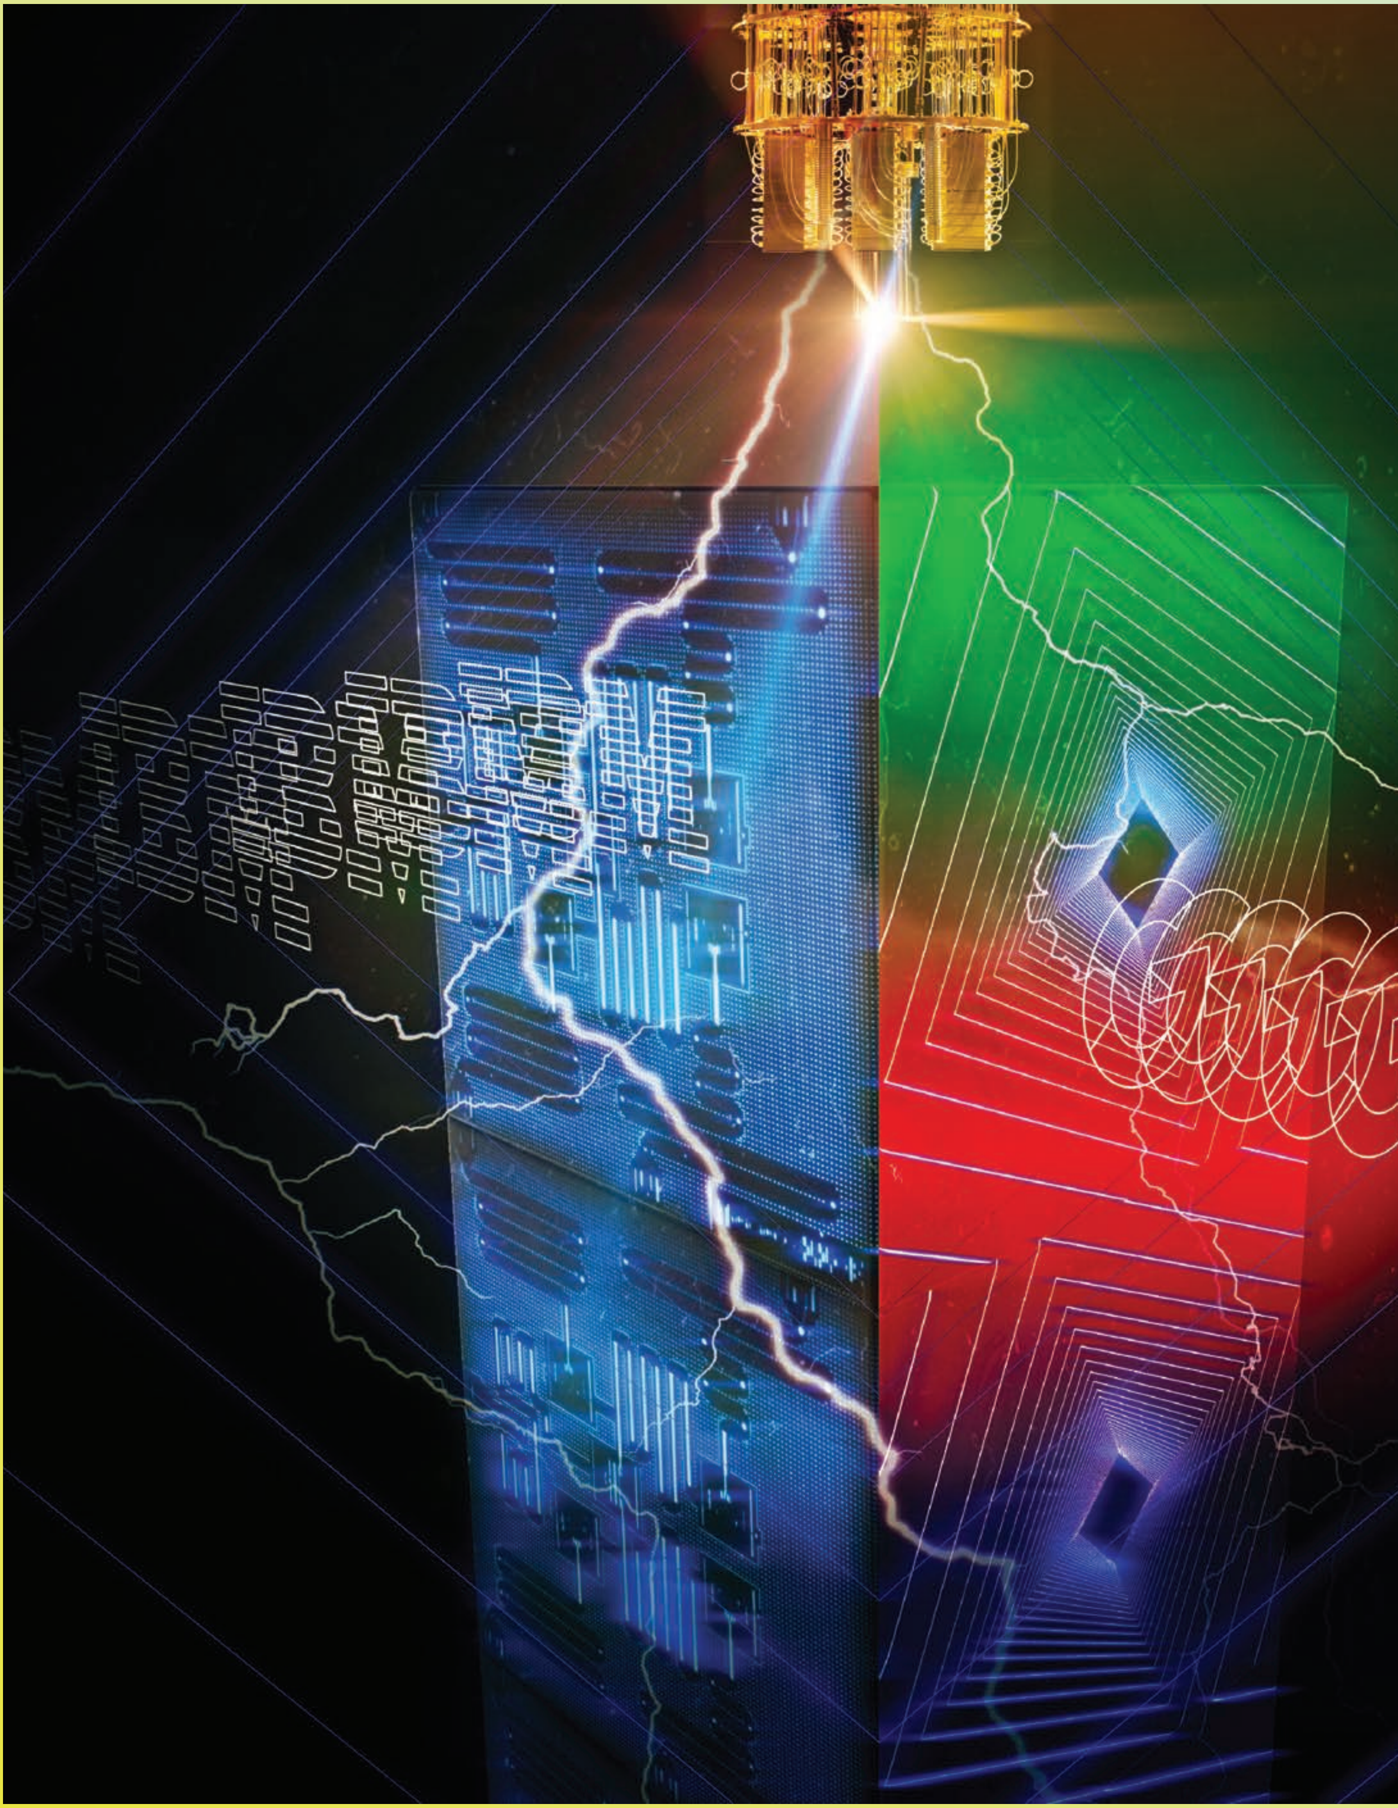

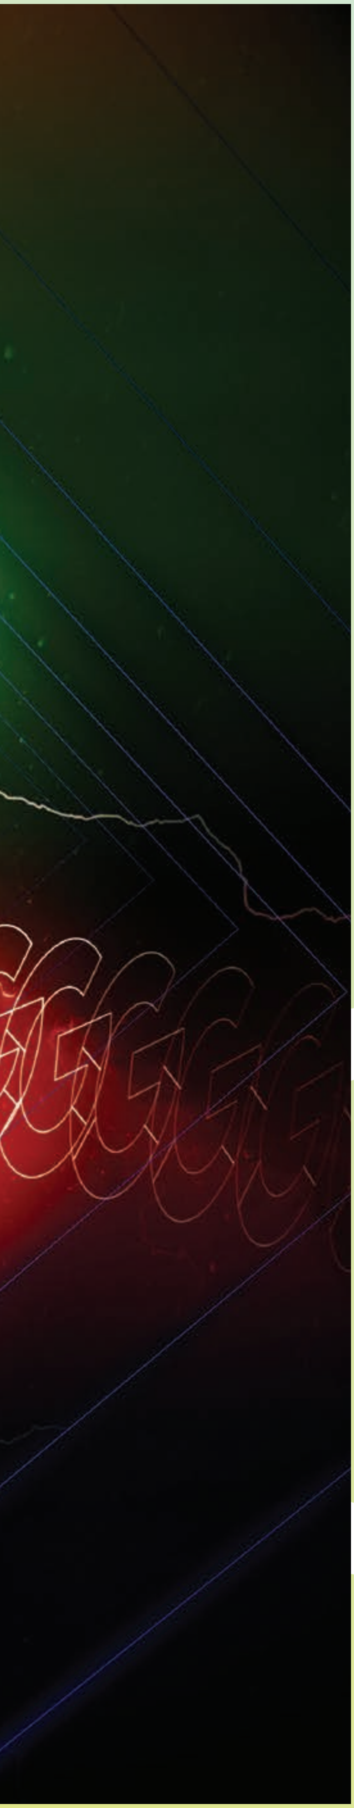

YOSHI SODEOKA

QUANTUM SUPREMACY

Google has provided the first clear proof of a quantum computer outperforming a classical one.

WHY IT MATTERS

Eventually, quantum computers will be able to solve problems no classical machine can manage.

KEY PLAYERS

Google

IBM

Microsoft

Rigetti

D-Wave

IonQ

Zapata Computing

Quantum Circuits

AVAILABILITY

5-10+ years

Quantum computers store and process data in a way completely differently from the ones we're all used to. In theory, they could tackle certain classes of problems that even the most powerful classical supercomputer imaginable would take millennia to solve, like breaking today's cryptographic codes or simulating the precise behavior of molecules to help discover new drugs and materials.

There have been working quantum computers for several years, but it's only under certain conditions that they outperform classical ones, and in October Google claimed the first such demonstration of "quantum supremacy." A computer with 53 qubits—the basic unit of quantum computation—did a calculation in a little over three minutes that, by Google's reckoning, would have taken the world's biggest supercomputer 10,000 years, or 1.5 billion times as long. IBM challenged Google's claim, saying the speedup would be a thousandfold at best; even so, it was a milestone, and each additional qubit will make the computer twice as fast.

However, Google's demo was strictly a proof of concept—the equivalent of doing random sums on a calculator and showing that the answers are right. The goal now is to build machines with enough qubits to solve useful problems. This is a formidable challenge: the more qubits you have, the harder it is to maintain their delicate quantum state. Google's engineers believe the approach they're using can get them to somewhere between 100 and 1,000 qubits, which may be enough to do something useful—but nobody is quite sure what.

And beyond that? Machines that can crack today's cryptography will require millions of qubits; it will probably take decades to get there. But one that can model molecules should be easier to build.

—Gideon Lichfield

We can now run powerful AI algorithms on our phones.

TINY AI

1

WHY IT MATTERS

Our devices no longer need to talk to the cloud for us to benefit from the latest AI-driven features.

KEY PLAYERS

Google

IBM

Apple

Amazon

AVAILABILITY

Now

AI has a problem: in the quest to build more powerful algorithms, researchers are using ever greater amounts of data and computing power, and relying on centralized cloud services. This not only generates alarming amounts of carbon emissions but also limits the speed and privacy of AI applications.

But a countertrend of tiny AI is changing that. Tech giants and academic researchers are working on new algorithms to shrink existing deep-learning models without losing their capabilities. Meanwhile, an emerging generation of specialized AI chips promises to pack more computational power into tighter physical spaces, and train and run AI on far less energy.

These advances are just starting to become available to consumers. Last May, Google announced that it can now run Google Assistant on users' phones without sending requests to a remote server. As of iOS 13, Apple runs Siri's speech recognition capabilities and its QuickType keyboard locally on the iPhone. IBM and Amazon now also offer developer platforms for making and deploying tiny AI.

All this could bring about many benefits. Existing services like voice assistants, autocorrect, and digital cameras will get better and faster without having to ping the cloud every time they need access to a deep-learning model. Tiny AI will also make new applications possible, like mobile-based medical-image analysis or self-driving cars with faster reaction times. Finally, localized AI is better for privacy, since your data no longer needs to leave your device to improve a service or a feature.

But as the benefits of AI become distributed, so will all its challenges. It could become harder to combat surveillance systems or deepfake videos, for example, and discriminatory algorithms could also proliferate. Researchers, engineers, and policymakers need to work together now to develop technical and policy checks on these potential harms. —*Karen Hao*

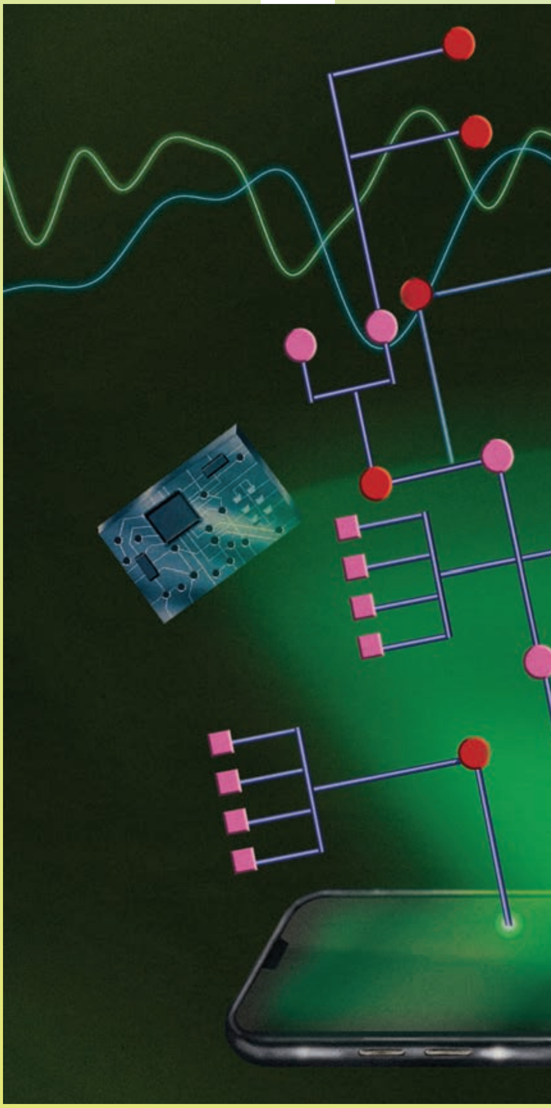

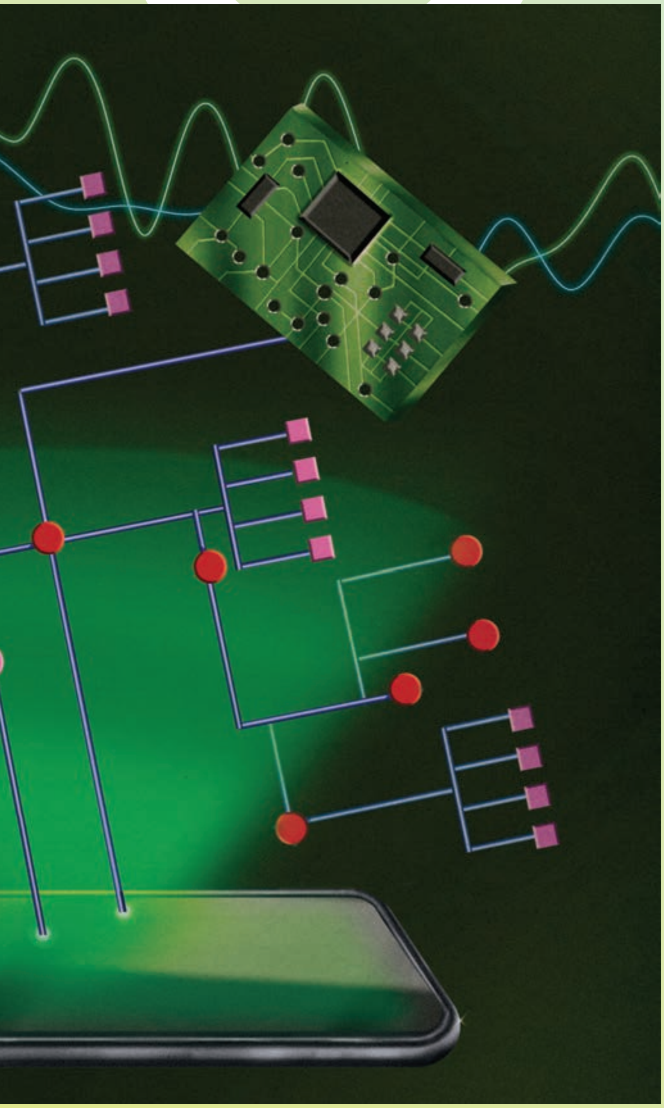

JULIA DUFOSSÉ

DIFFERENTIAL PRIVACY

A technique to measure the privacy of a crucial data set.

WHY IT MATTERS

It is increasingly difficult for the US Census Bureau to keep the data it collects private. A technique called differential privacy could solve that problem, build trust, and also become a model for other countries.

KEY PLAYERS

US Census Bureau

Apple

Facebook

AVAILABILITY

Its use in the 2020 US Census will be the biggest-scale application yet.

In 2020, the US government has a big task: collect data on the country's 330 million residents while keeping their identities private. The data is released in statistical tables that policymakers and academics analyze when writing legislation or conducting research. By law, the Census Bureau must make sure that it can't lead back to any individuals.

But there are tricks to "de-anonymize" individuals, especially if the census data is combined with other public statistics.

So the Census Bureau injects inaccuracies, or "noise," into the data. It might make some people younger and others older, or label some white people as black and vice versa, while keeping the totals of each age or ethnic group the same. The more noise you inject, the harder de-anonymization becomes.

Differential privacy is a mathematical technique that makes this process rigorous by measuring how much privacy increases when noise is added. The method is already used by Apple and Facebook to collect aggregate data without identifying particular users.

But too much noise can render the data useless. One analysis showed that a differentially private version of the 2010 Census included households that supposedly had 90 people.

If all goes well, the method will likely be used by other federal agencies. Countries like Canada and the UK are watching too. —*Angela Chen*

Researchers can now spot climate change's role in extreme weather.

CLIMATE CHANGE ATTRIBUTION

Ten days after Tropical Storm Imelda began flooding neighborhoods across the Houston area last September, a rapid-response research team announced that climate change almost certainly played a role.

The group, World Weather Attribution, had compared high-resolution computer simulations of worlds where climate change did and didn't occur. In the former, the world we live in, the severe storm was as much as 2.6 times more likely—and up to 28% more intense.

Earlier this decade, scientists were reluctant to link any specific event to climate change. But many more extreme-weather attribution studies have been done in the last few years, and rapidly improving tools and techniques have made them more reliable and convincing.

This has been made possible by a combination of advances. For one, the lengthening record of detailed satellite data is helping us understand natural systems. Also, increased computing power means scientists can create higher-resolution simulations and conduct many more virtual experiments.

These and other improvements have allowed scientists to state with increasing statistical certainty that yes, global warming is often fueling more dangerous weather events.

By disentangling the role of climate change from other factors, the studies are telling us what kinds of risks we need to prepare for, including how much flooding to expect and how severe heat waves will get as global warming becomes worse. If we choose to listen, they can help us understand how to rebuild our cities and infrastructure for a climate-changed world. —*James Temple*

WHY IT MATTERS

It's providing a clearer sense of how climate change is worsening the weather, and what we'll need to do to prepare.

KEY PLAYERS

World Weather Attribution

Royal Netherlands Meteorological Institute

Red Cross Red Crescent Climate Centre

AVAILABILITY

Now

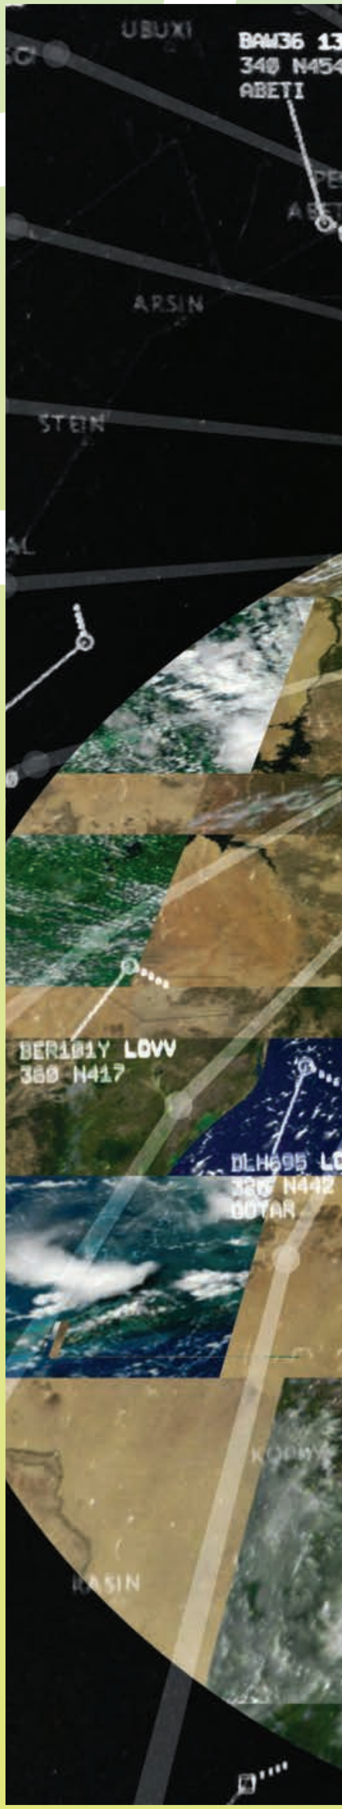

YOSHI SODEOKA


320 N440
BEGLA
BAW146 133,985
360 N423
ABETI
VEDER
VAMOC
XOMBA
ALAMU
ERCOM
ASL74V
190 N271
BABIT
AUA829 E3U
270 ↑ N417
360 N415
AUA783R
260 ↑ N424
310 370
SIRDU
DLH1298 E3U
350 N474
INVED
THY4CV W2M
380 N415
GOTAR
THY6QL W2M
380 VIR N416
380 360 340
SUNIS
AUA279M LYBA
320 ↑ N419
330
BABIT
BAW664 LYBA
350 N470
LZB301 LOVV
380 N418
BEGLA
SWR146W
270 ↑ N453
340 380
GOTAR
BAW670 W2M
TONDO

SKY'S THE LIMIT

How mega-constellations will change the way we use space.

By RAMIN SKIBBA

Sixty-three years after *Sputnik* first entered orbit, a couple of thousand satellites circle the planet to help us do things like communicate, navigate, and forecast the weather. Soon, though, they will be dwarfed by mega-constellations with great networks of hundreds or even thousands of satellites working in concert.

Starlink, from Elon Musk's SpaceX, offers the clearest glimpse of what's to come. The company has already deployed more than 100 satellites for the system, and by the mid-2020s, it plans to assemble a constellation of nearly 12,000 to provide broadband internet access globally. Many other space agencies and for-profit space companies have begun setting up their own networks, too.

"It's a rather dynamic environment right now, with a lot of people starting to look at space as a means to answer certain business models," says Roger Hunter, manager of NASA's Small Spacecraft Technology program. "I call it the democratization of space."

Constellations offer new levels of versatility. Smaller, cheaper satellites—some just the size of a briefcase—can be arranged in different configurations depending on their goal. Lined up in a string that follows a single orbit, for example, a constellation can repeatedly photograph or surveil the same spot. Starlink, meanwhile, is arranged in a crisscross formation to blanket the planet with internet service.

"I think that as an industry we're trying to figure out how to increase the level of great space-based services that come down and help people on Earth every day, while doing it in a responsible and sustainable way in the orbital environment," says Mike Safyan, vice president of launch and global ground systems at Planet Labs, which operates the second-largest constellation in operation.

In the meantime, we can look forward to more and bigger satellite systems, with hundreds if not thousands of members, heading up into orbit. And eventually, wherever humans go—whether it's to the moon, Mars, or even other stars—they'll be taking constellations with them. ■

Ramin Skibba is an astrophysicist turned science writer.

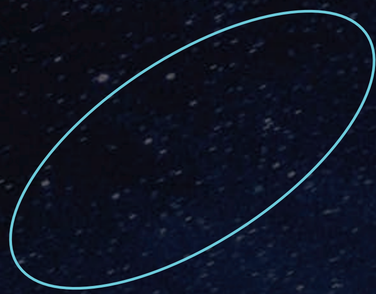

SINGLE ORBITAL PLANES CAN REPEATEDLY VISIT THE SAME SPOT MULTIPLE TIMES A DAY.

GRID ORBITS ARE LAYERED IN ORDER TO GIVE EVEN GREATER COVERAGE—GREAT FOR PROVIDING INTERNET ACCESS.

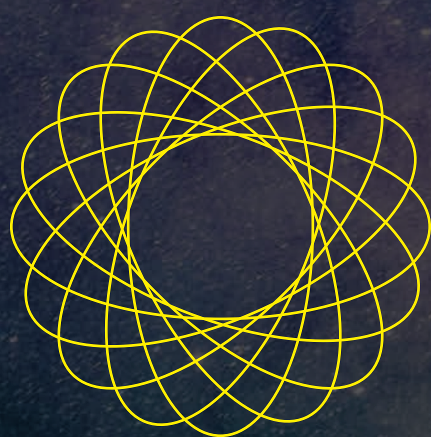

CRISSCROSSING ORBITS ARE TYPICALLY USED FOR COMMUNICATIONS OR GUIDANCE AND NAVIGATION.

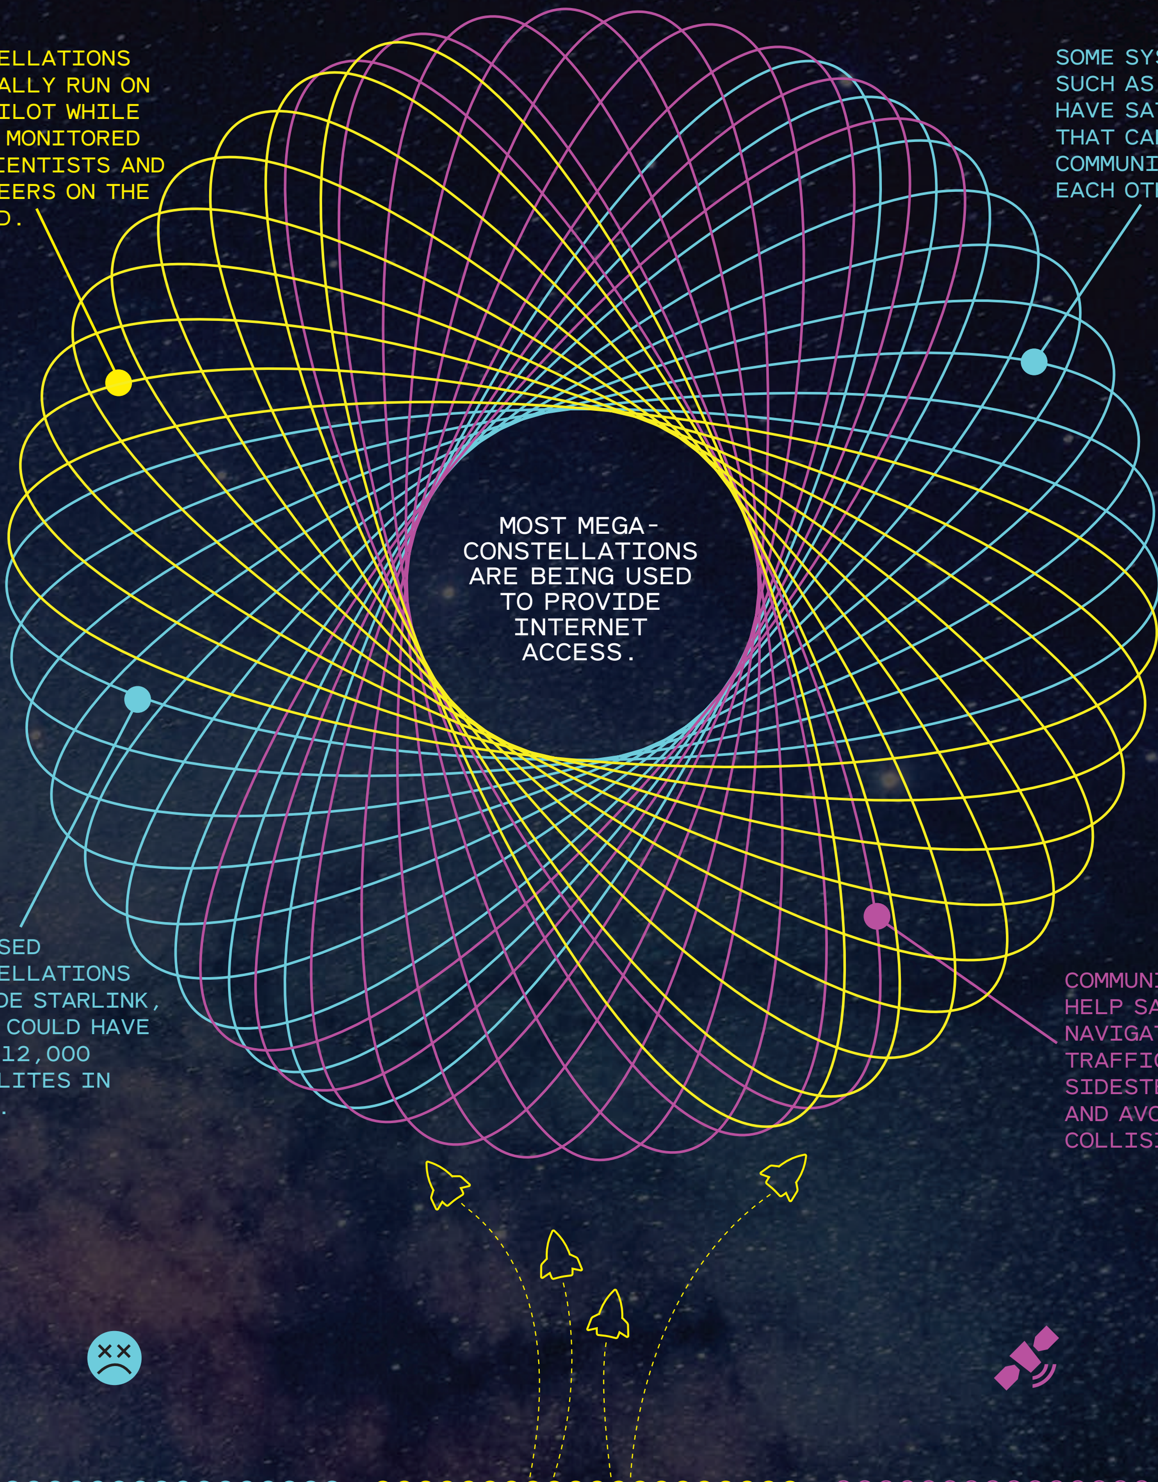

RESILIENCE

If one satellite fails, others can step in to cover.

Substitutions make sure the system keeps going if a single unit breaks.

Dying satellites get dragged into a low orbit and burn up on reentry.

LAUNCH

A single rocket carries up to 60 satellites at a time.

Batched launches mean a whole operation won't be lost if a rocket fails.

More satellites can join the formation later.

EQUIPMENT

Cubesats are commonly used; they are shoebox-like and weigh only 4 to 5 kilograms.

Planet Labs's SkySats are the size of a mini-fridge and weigh 100 kg.

The company's entire fleet weighs half as much as one ordinary high-resolution imaging satellite.

CASH IS GRADUALLY DYING OUT. WILL WE EVER HAVE A DIGITAL ALTERNATIVE THAT OFFERS THE

SAME MIX OF CONVENIENCE AND FREEDOM?

AN ELEGY FOR CASH

THE TECHNOLOGY WE MIGHT NEVER REPLACE

By MIKE ORCUTT / Illustrations by ANDREA DAQUINO

Think about the last time you used cash. How much did you spend? What did you buy, and from whom? Was it a one-time thing, or was it something you buy regularly?

Was it legal?

If you'd rather keep all that to yourself, you're in luck. The person in the store (or on the street corner) may remember your face, but as long as you didn't reveal any identifying information, there is nothing that links you to the transaction.

This is a feature of physical cash that payment cards and apps do not have: freedom. Called "bearer instruments," banknotes and coins are presumed to be owned by whoever holds them. We can use them to transact with another person without a third party getting in the way. Companies cannot build advertising profiles or credit ratings out of our data, and governments cannot track our spending or our movements. And while a credit card can be declined and a check mislaid, handing over money works every time, instantly.

We shouldn't take this freedom for granted. Much of our commerce now happens online. It relies on banks and financial technology companies to serve as middlemen. Transactions are going digital in the physical world, too: electronic payment tools, from debit cards to Apple Pay to Alipay, are increasingly replacing cash. While notes and coins remain popular in many countries, including the US, Japan, and Germany, in others they are nearing obsolescence.

This trend has civil liberties groups worried. Without cash, there is "no chance for the kind of dignity-preserving privacy that undergirds an open society," writes Jerry Brito, executive director of Coin Center, a policy advocacy group based in Washington, DC. In a recent report, Brito says we must "develop and foster" electronic cash that is as private as physical cash and doesn't require permission to use.

The central question is who will develop and control the electronic payment systems of the future. Most of the existing ones, like Alipay, Zelle, PayPal, Venmo, and Kenya's M-Pesa, are run by private firms. Afraid of leaving payments solely in their hands, many governments are looking to develop some sort of electronic stand-in for notes and coins. Meanwhile, advocates of stateless, ownerless cryptocurrencies like Bitcoin say they're the only solution as surveillance-proof as cash—but can they be feasible at large scales?

We tend to take it for granted that new technologies work better than old ones—safer, faster, more accurate, more efficient, more convenient. Purists may extol the virtues of vinyl records, but nobody can dispute that a digital music collection is easier to carry and sounds almost exactly as good. Cash is a paradox—a technology thousands of years old that may just prove impossible to re-create in a more advanced form.

IN (GOVERNMENT) MONEY WE TRUST?

We call banknotes and coins "cash," but the term really refers to something more abstract: cash is essentially money that your government owes you. In the old days this was a literal debt. "I promise to pay the bearer on demand the sum of…" still appears on British banknotes, a notional guarantee that the Bank of England will hand over the same value in gold in exchange for your note. Today it represents the more abstract guarantee that you will always be able to use that note to pay for things.

The digits in your bank account, on the other hand, refer to what your bank owes you. When you go to an ATM, you are effectively converting the bank's promise to pay into a government promise.

Most people would say they trust the government's promise more, says Gabriel Söderberg, an economist at the Riksbank, the central bank of Sweden. Their bet—correct, in most countries—is that their government is much less likely to go bust.

That's why it would be a problem if Sweden were to go completely "cashless," Söderberg says. He and his colleagues fear that if people lose the option to convert their bank money to government money at will and use it to pay for whatever they need, they might start to lose trust in the whole money system. A further worry is that if the private sector is left to dominate digital payments, people who can't or won't use these systems could be shut out of the economy.

This is fast becoming more than just a thought experiment in Sweden. Nearly everyone there uses a mobile app called Swish to pay for things. Economists have estimated that retailers in Sweden could completely stop accepting cash by 2023.

Creating an electronic version of Sweden's sovereign currency—an "e-krona"—could mitigate these problems, Söderberg says. If the central bank were to issue digital money, it would design it to be a public good, not a profit-making product for a corporation. "Easily accessible, simple, and user-friendly versions could be developed for those who currently have difficulty with digital technology," the bank asserted in a November report covering Sweden's payment landscape.

The Riksbank plans to develop and test an e-krona prototype. It has examined a number of technologies that might underlie it, including cryptocurrency systems like Bitcoin. But the central bank has also called on the Swedish government to lead a broad

THE DECLINE OF CASH

% of transactions (by value) in cash for selected countries*

*ACCORDING TO IMF'S CASHSHARE METRIC; SOURCE: INTERNATIONAL MONETARY FUND

public inquiry into whether such a system should ever go live. "In the end, this decision is too big for a central bank alone, at least in the Swedish context," Söderberg says.

THE DEATH OF FINANCIAL PRIVACY

China, meanwhile, appears to have made its decision: the digital renminbi is coming. Mu Changchun, head of the People's Bank of China's digital currency research institute, said in September that the currency, which the bank has been working on for years, is "close to being out." In December, a local news report suggested that the PBOC is nearly ready to start tests in the cities of Shenzhen and Suzhou. And the bank has been explicit about its intention to use it to replace banknotes and coins.

Cash is already dying out on its own in China, thanks to Alipay and WeChat Pay, the QR-code-based apps that have become ubiquitous in just a few years. It's been estimated that mobile payments made up more than 80% of all payments in China in 2018, up from less than 20% in 2013.

It's not clear how much access the government currently has to transaction data from WeChat Pay and Alipay. Once it issues a sovereign digital currency—which officials say will be compatible with those two services—it will likely have access to a lot more. Martin Chorzempa, a research fellow at the Peterson Institute for International Economics in Washington, DC, told the New York Times in October that the system will give the PBOC "extraordinary power and visibility into the financial system, more than any central bank has today."

We don't know for sure what technology the PBOC plans to use as the basis for its digital renminbi, but we have at least two revealing clues. First, the bank has been researching blockchain technology since 2014, and the government has called the development of this technology a priority. Second, Mu said in September that China's system will bear similarities to Libra, the electronic currency Facebook announced last June. Indeed, PBOC officials have implied in public statements that the unveiling of Libra inspired them to accelerate the development of the digital renminbi, which has been in the works for years.

As currently envisioned, Libra will run on a blockchain, a type of accounting ledger that can be maintained by a network of computers instead of a single central authority. However, it will operate very differently from Bitcoin, the original blockchain system.

The computers in Bitcoin's network use open-source software to automatically verify and record every single transaction. In the process, they generate a permanent public record of the currency's entire transaction history: the blockchain. As envisioned, Libra's network will do something similar. But whereas anyone with a computer and an internet connection can participate anonymously in Bitcoin's network, the "nodes" that make up Libra's network will be companies that have been vetted and given membership in a nonprofit association.

Unlike Bitcoin, which is notoriously volatile, Libra will be designed to maintain a stable value. To this end, the so-called Libra Association will be responsible for maintaining a reserve of government-issued currencies (the latest plan is for it to be half US dollars, with the other half composed of British pounds, euros, Japanese yen, and Singapore dollars). This reserve is supposed to serve as backing for the digital units of value.

Both Libra and the digital renminbi, however, face serious questions about privacy. To start with, it's not clear if people will be able to use them anonymously.

With Bitcoin, although transactions are public, users don't have to reveal who they really are; each person's "address" on the public blockchain is just a random string of letters and numbers. But in recent years, law enforcement officials have grown skilled at combining public blockchain data with other clues to unmask people using cryptocurrencies for illicit purposes. Indeed, in a July blog post, Libra project head David Marcus argued that the currency would be a boon for law enforcement, since it would help "move more cash transactions—where a lot of illicit activities happen—to a digital network."

As for the Chinese digital currency, Mu has said it will feature some level of anonymity. "We know the demand from the general public is to keep anonymity by using paper money and coins ... we will give those people who demand it anonymity," he said at a November conference in Singapore. "But at the same time we will keep the balance between 'controllable anonymity' and anti-money-laundering, CTF [counter-terrorist financing], and

also tax issues, online gambling, and any electronic criminal activities," he added. He did not, however, explain how that "balance" would work.

Sweden and China are leading the charge to issue consumer-focused electronic money, but according to the International Monetary Fund, more than 20 countries appear to be at least exploring the idea. In some, the rationale is similar to Sweden's: dwindling cash and a growing private-sector payments ecosystem. Others are countries where commercial banks have decided not to set up shop. Many see an opportunity to better monitor for illicit transactions. All will have to wrestle with the same thorny privacy issues that Libra and the digital renminbi are raising.

Robleh Ali, a research scientist at MIT's Digital Currency Initiative, says digital currency systems from central banks may need to be designed so that the government can "consciously blind itself" to the information. Something like that might be technically possible thanks to cutting-edge cryptographic tools like zero-knowledge proofs, which are used in systems like Zcash to shield blockchain transaction information from public view.

However, there's no evidence that any governments are even thinking about deploying tools like this. And regardless, can any government—even Sweden's—really be trusted to blind itself?

CRYPTOCURRENCY: A WORKAROUND FOR FREEDOM

That's wishful thinking, says Alex Gladstein, chief strategy officer for the Human Rights Foundation. While you may trust your government or think you've got nothing to hide, that might not always remain true. Politics evolves, governments get pushed out by elections or other events, what constitutes a "crime" changes, and civil liberties are not guaranteed. "Financial privacy is not going to be gifted to you by your government, regardless of how 'free' they are," Gladstein says. He's convinced that it has to come in the form of a stateless, decentralized digital currency like Bitcoin.

In fact, "electronic cash" was what Bitcoin's still-unknown inventor, the pseudonymous Satoshi Nakamoto, claimed to be trying to create (before disappearing). Eleven years into its life, Nakamoto's technology still lacks some of the signature features of cash. It is difficult to use, transactions can take more than an hour to process, and the currency's value can fluctuate wildly. And as already noted, the supposedly anonymous transactions it enables can sometimes be traced.

But in some places people just need something that works, however imperfectly. Take Venezuela. Cash in the crisis-ridden country is scarce, and the Venezuelan bolivar is constantly losing value to hyperinflation. Many Venezuelans seek refuge in US dollars, storing them under the proverbial (and literal) mattress, but that also makes them vulnerable to thieves.

What many people want is access to stable cash in digital form, and there's no easy way to get that, says Alejandro Machado, cofounder of the Open Money Initiative. Owing to government-imposed capital controls, Venezuelan banks have largely been cut off from foreign banks. And under restrictions by US financial institutions, digital money services like PayPal and Zelle are inaccessible to most people. So a small number of tech-savvy Venezuelans have turned to a service called LocalBitcoins.

It's like Craigslist, except that the only things for sale are bitcoins and bolivars. On Venezuela's LocalBitcoins site, people advertise varying quantities of currency for sale at varying exchange rates. The site holds the money in escrow until trades are complete, and tracks the sellers' reputations.

It's not for the masses, but it's "very effective" for people who can make it work, says Machado. For instance, he and his colleagues met a young woman who mines Bitcoin and keeps her savings in the currency. She doesn't have a foreign bank account, so she's willing to deal with the constant fluctuations in Bitcoin's price. Using LocalBitcoins, she can cash out into bolivars whenever she needs them—to buy groceries, for example. "Niche power users" like this are "leveraging the best features of Bitcoin, which is to be an asset that is permissionless and that is very easy to trade electronically," Machado says.

However, this is possible only because there are enough people using the site to create what finance people call "local liquidity," meaning you can easily find a buyer for your bitcoins or bolivars. Bitcoin is the only cryptocurrency that has achieved this in Venezuela, says Machado, and it's mostly thanks to LocalBitcoins.

This is a long way from the dream of cryptocurrency as a widely used substitute for stable, government-issued money. Most Venezuelans can't use Bitcoin, and few merchants there even know what it is, much less how to accept it.

Still, it's a glimpse of what a cryptocurrency can offer—a functional financial system that anyone can join and that offers the kind of freedom cash provides in most other places.

DECENTRALIZE THIS

Could something like Bitcoin ever be as easy to use and reliable as today's cash is for everyone else? The answer is philosophical as well as technical.

To begin with, what does it even mean for something to be "like Bitcoin"? Central banks and corporations will adapt certain aspects of Bitcoin and apply them to their own ends. Will those be cryptocurrencies? Not according to purists, who say that though Libra or some future central-bank-issued digital currency may run on blockchain technology, these won't be cryptocurrencies because they will be under centralized control.

HOW BIG A PROBLEM IS THIS? THAT DEPENDS ON WHERE YOU LIVE, HOW MUCH YOU TRUST YOUR GOVERNMENT AND YOUR FELLOW CITIZENS, AND WHY YOU WISH TO USE CASH.

True cryptocurrencies are "decentralized"—they have no one entity in charge and no single points of failure, no weak spots that an adversary (including a government) could attack. With no middleman like a bank attesting that a transaction took place, each transaction has to be validated by the nodes in a cryptocurrency's network, which can number many thousands. But this requires an immense expenditure of computing power, and it's the reason Bitcoin transactions can take more than an hour to settle.

A currency like Libra wouldn't have this problem, because only a few authorized entities would be able to operate nodes. The trade-off is that its users wouldn't be able to trust those entities to guarantee their privacy, any more than they can trust a bank, a government, or Facebook.

Is it technically possible to achieve Bitcoin's level of decentralization and the speed, scale, privacy, and ease of use that we've come to expect from traditional payment methods? That's a problem many talented researchers are still trying to crack. But some would argue that shouldn't necessarily be the goal.

In a recent essay, Jill Carlson, cofounder of the Open Money Initiative, argued that decentralized cryptocurrency systems "were not meant to go mainstream." Rather, they were created explicitly for "censored transactions," from paying for drugs or sex to supporting political dissidents or getting money out of countries with restrictive currency controls. Their slowness is inherent, not a design flaw; they "forsake scale, speed, and cost in favor of one key feature: censorship resistance." A world in which they went mainstream would be "a very scary place indeed," she wrote.

In summary, we have three avenues for the future of digital money, none of which offers the same mix of freedom and ease of use that characterizes cash. Private companies have an obvious incentive to monetize our data and pursue profits over public interest. Digital government money may still be used to track us, even by well-intentioned governments, and for less benign ones it's a fantastic tool for surveillance. And cryptocurrency can prove useful when freedoms are at risk, but it likely won't work at scale anytime soon, if ever.

How big a problem is this? That depends on where you live, how much you trust your government and your fellow citizens, and why you wish to use cash. And if you'd rather keep that to yourself, you're in luck. For now. ■

Mike Orcutt is MIT Technology Review's senior blockchain reporter.

TO REIGN SUPREME

BY GIDEON LICHFIELD

WHAT THE RACE TO BUILD A QUANTUM COMPUTER REVEALS ABOUT GOOGLE AND IBM.

RIGETTI

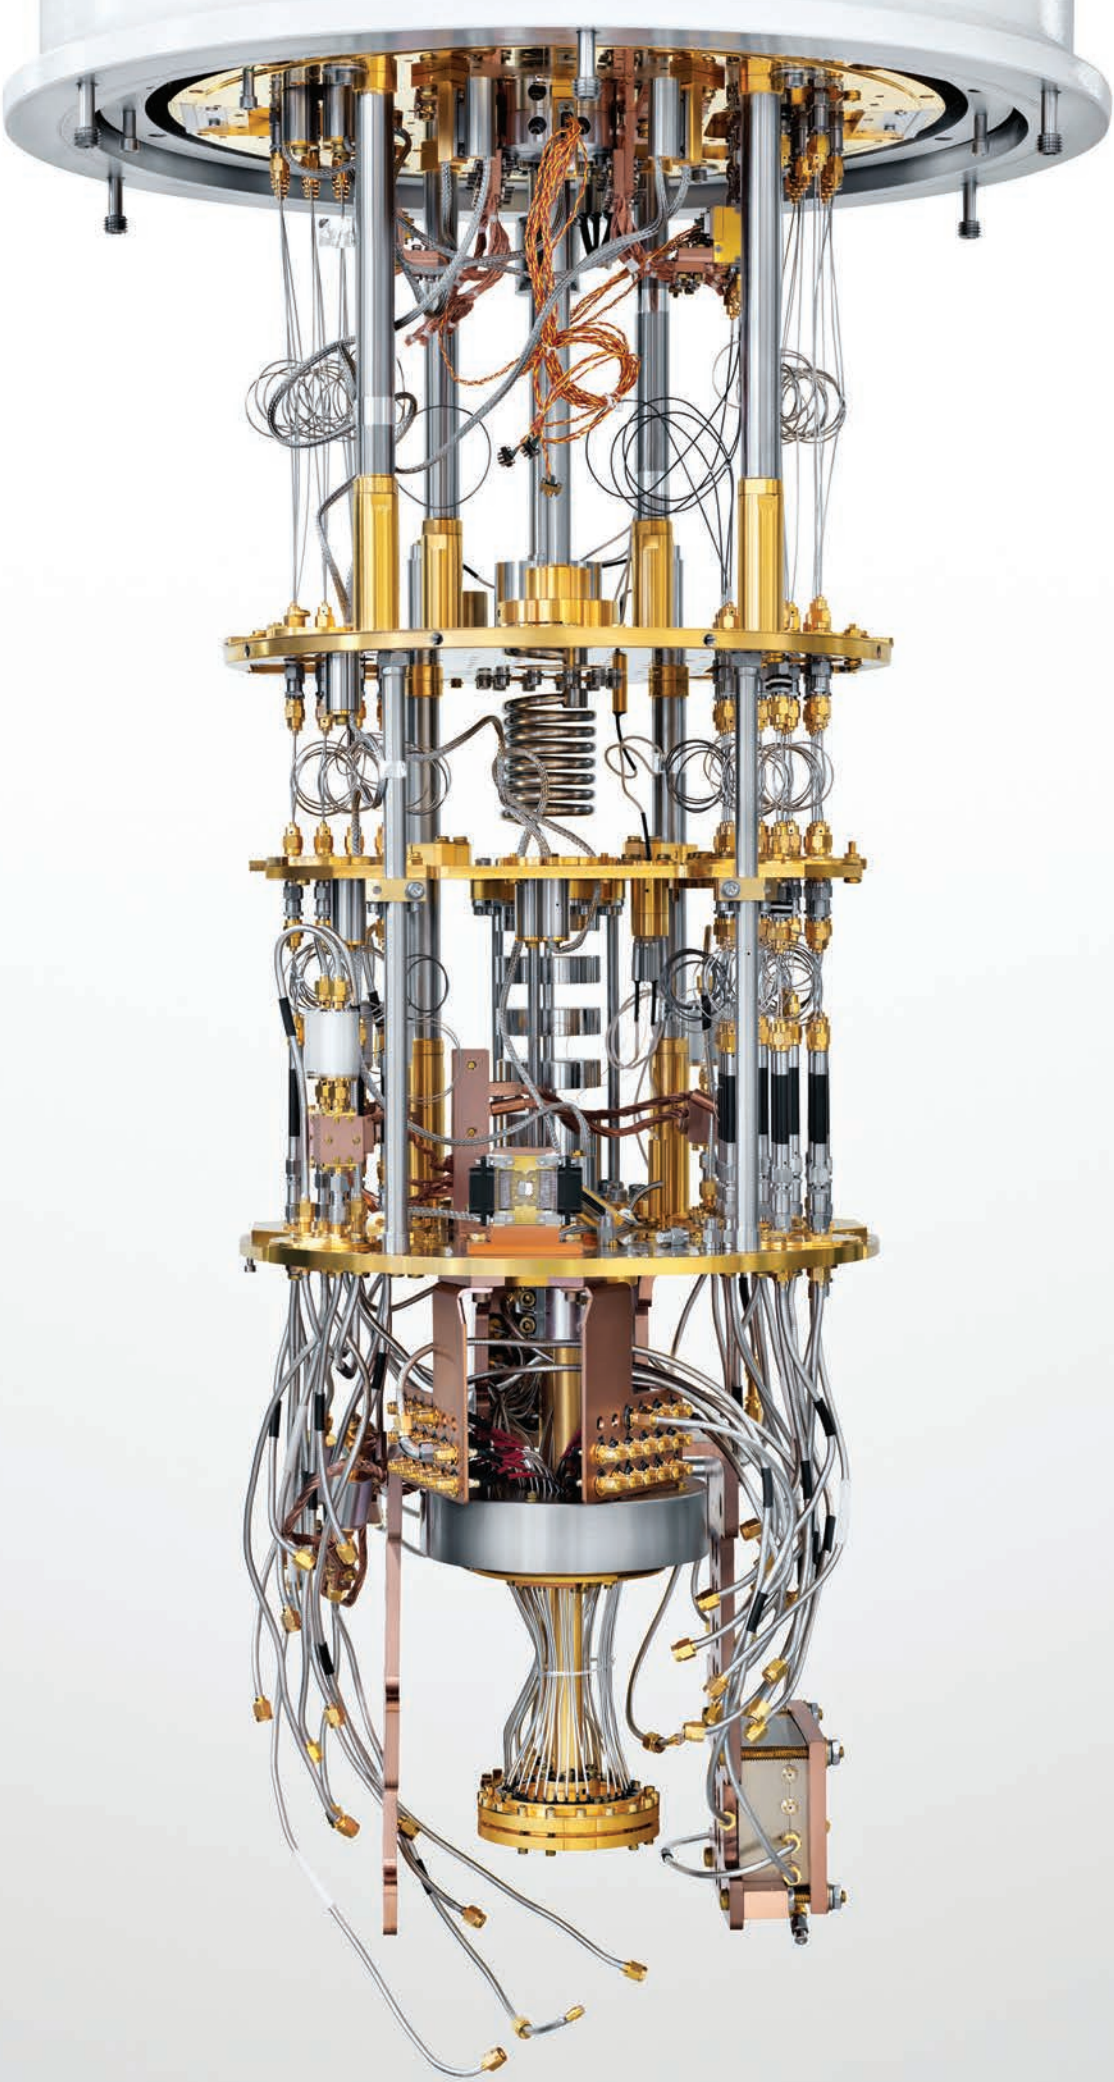

Google's most advanced computer isn't at the company's headquarters in Mountain View, California, nor anywhere in the febrile sprawl of Silicon Valley. It's a few hours' drive south in Santa Barbara, in a flat, soulless office park inhabited mostly by technology firms you've never heard of.

An open-plan office holds several dozen desks. There's an indoor bicycle rack and designated "surfboard parking," with boards resting on brackets that jut out from the wall. Wide double doors lead into a lab the size of a large classroom. There, amidst computer racks and jumbles of instrumentation, a handful of cylindrical vessels—each a little bigger than an oil drum—hang from vibration-damping rigs like enormous steel pupae.

On one of them, the outer vessel has been removed to expose a multi-tiered tangle of steel and brass innards known as "the chandelier." It's basically a supercharged refrigerator that gets colder with each layer down. At the bottom, kept in a vacuum a hair's breadth above absolute zero, is what looks to the naked eye like an ordinary silicon chip. But rather than transistors, it's etched with tiny superconducting circuits that, at these low temperatures, behave as if they were single atoms obeying the laws of quantum physics. Each one is a quantum bit, or qubit—the basic information-storage unit of a quantum computer.

Late last October, Google announced that one of those chips, called Sycamore, had become the first to demonstrate "quantum supremacy" by performing a task that would be practically impossible on a classical machine. With just 53 qubits, Sycamore had completed a calculation in a few minutes that, according to Google, would have taken the world's most powerful existing supercomputer, Summit, 10,000 years. Google touted this as a major breakthrough, comparing it to the launch of *Sputnik* or the first flight by the Wright brothers—the threshold of a new era of machines that would make today's mightiest computer look like an abacus.

At a press conference in the lab in Santa Barbara, the Google team cheerfully fielded questions from journalists for nearly three hours. But their good humor couldn't quite mask an underlying tension. Two days earlier, researchers from IBM, Google's leading rival in quantum computing, had torpedoed its big reveal. They'd published a paper that essentially accused the Googlers of getting their sums wrong. IBM reckoned it would have taken Summit merely days, not millennia, to replicate what Sycamore had done. When asked what he thought of IBM's result, Hartmut Neven, the head of the Google team, pointedly avoided giving a direct answer.

You could dismiss this as just an academic spat—and in a sense it was. Even if IBM was right, Sycamore had still done the calculation a thousand times faster than Summit would have. And it would likely be only months before Google built a slightly larger quantum machine that proved the point beyond doubt.

IBM's deeper objection, though, was not that Google's experiment was less successful than claimed, but that it was a meaningless test in the first place. Unlike most of the quantum computing world, IBM doesn't think "quantum supremacy" is the technology's Wright brothers moment; in fact, it doesn't even believe there will be such a moment.

IBM is instead chasing a very different measure of success, something it calls "quantum advantage." This isn't a mere difference of words or even of science, but a philosophical stance with roots in IBM's history, culture, and ambitions—and, perhaps, the fact that for eight years its revenue and profit have been in almost unremitting decline, while Google and its parent company Alphabet have only seen their numbers grow. This context, and these differing goals, could influence which—if either—comes out ahead in the quantum computing race.

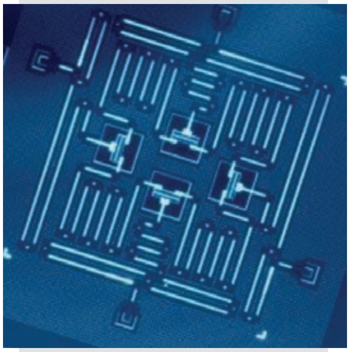

JAY M. GAMBETTA, JERRY M. CHOW & MATTHIAS STEFFEN

What's in a qubit?

Just as there were different transistor designs in the early days of computing, there are currently many ways to make qubits. Google and IBM both use a version of the leading method, a superconducting transmon qubit, of which the core component is a Josephson junction. This consists of a pair of superconducting metal strips separated by a gap just a nanometer wide; the quantum effects are a result of how electrons cross that gap.

Worlds apart

The sleek, sweeping curve of IBM's Thomas J. Watson Research Center in the suburbs north of New York City, a neo-futurist masterpiece by the Finnish architect Eero Saarinen, is a continent and a universe away from the Google team's nondescript digs. Completed in 1961 with the bonanza IBM made from mainframes, it has a museum-like quality, a reminder to everyone who works inside it of the company's breakthroughs in everything from fractal geometry to superconductors to artificial intelligence—and quantum computing.

The head of the 4,000-strong research division is Dario Gil, a Spaniard whose rapid-fire speech races to keep up with his almost evangelical zeal. Both times I spoke to him, he rattled off historical milestones intended to underscore how long IBM has been involved in quantum-computing-related research (see time line at right).

But over the decades, the company has gained a reputation for struggling to turn its research projects into commercial successes. Take, most recently, Watson, the *Jeopardy!*-playing AI that IBM tried to convert into a robot medical guru. It was meant to provide diagnoses and identify trends in oceans of medical data, but despite dozens of partnerships with health-care providers, there have been few commercial applications, and even the ones that did emerge have yielded mixed results.

The quantum computing team, in Gil's telling, is trying to break that cycle by doing the research and business development in parallel. Almost as soon as it had working quantum computers, it started making them accessible to outsiders by putting them on the cloud, where they can be programmed by means of a simple drag-and-drop interface that works in a web browser. The "IBM Q Experience," launched in 2016, now consists of 15 publicly available quantum computers ranging from five to 53 qubits in size. Some 12,000 people a month use them, ranging from academic researchers to school kids. Time on the smaller machines is free; IBM says it already has more than 100 clients paying (it won't say how much) to use the bigger ones.

None of these devices—or any other quantum computer in the world, except for Google's

How the quantum race stacks up so far

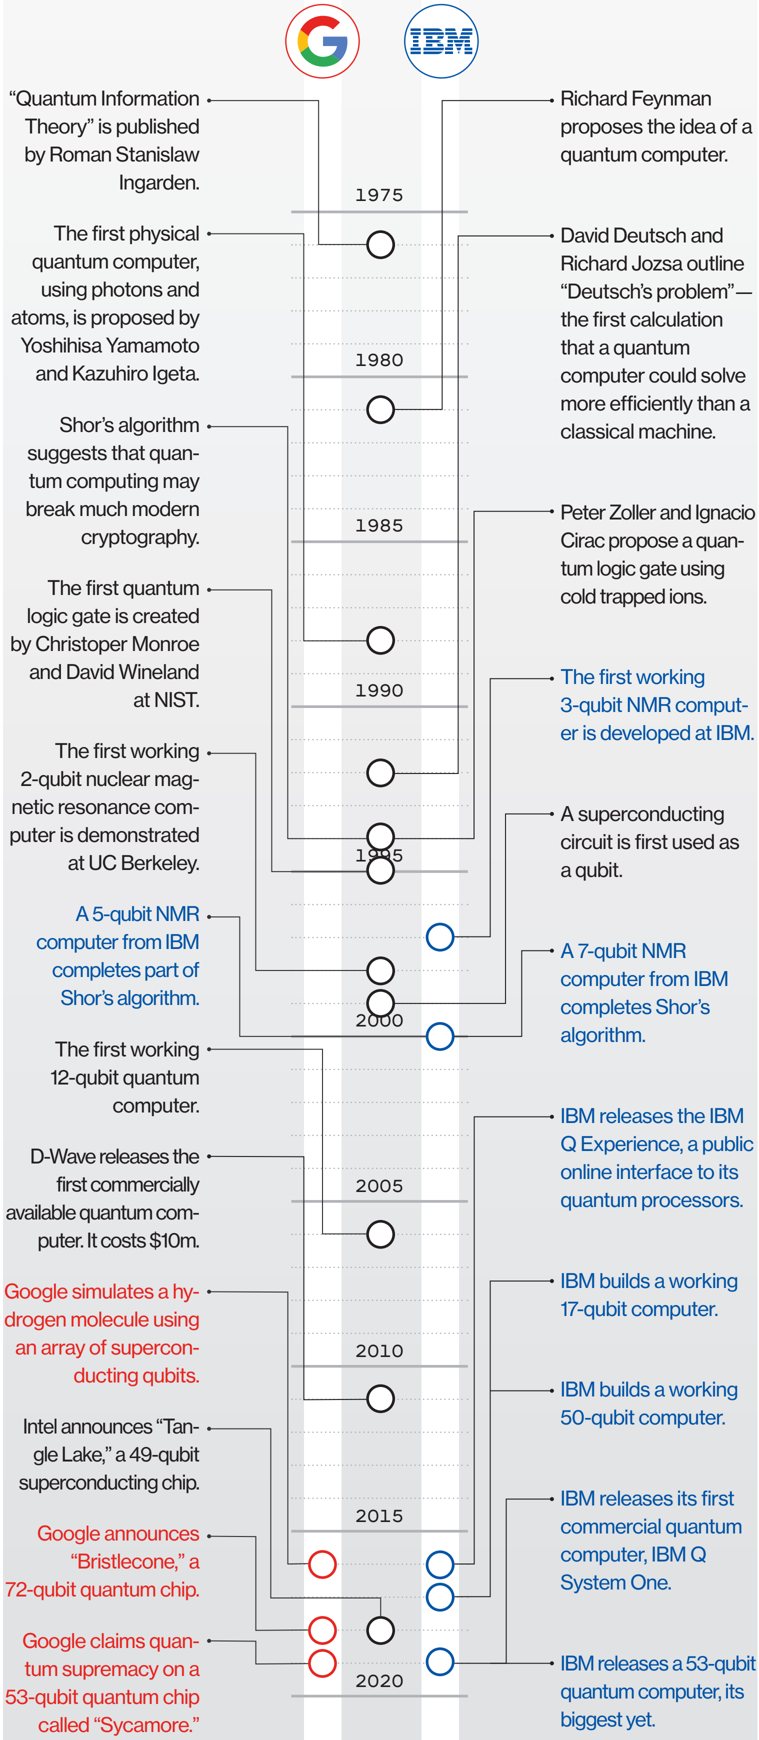

Sycamore—has yet shown it can beat a classical machine at anything. To IBM, that isn't the point right now. Making the machines available online lets the company learn what future clients might need from them and allows outside software developers to learn how to write code for them. That, in turn, contributes to their development, making subsequent quantum computers better.

This cycle, the company believes, is the fastest route to its so-called quantum advantage, a future in which quantum computers won't necessarily leave classical ones in the dust but will do *some* useful things *somewhat* faster or more efficiently—enough to make them economically worthwhile. Whereas quantum supremacy is a single milestone, quantum advantage is a "continuum," the IBMers say—a gradually expanding world of possibility.

This, then, is Gil's grand unified theory of IBM: that by combining its heritage, its technical expertise, other people's brainpower, and its dedication to business clients, it can build useful quantum computers sooner and better than anybody else.

In this view of things, IBM sees Google's quantum supremacy demonstration as "a parlor trick," says Scott Aaronson, a physicist at the University of Texas at Austin, who contributed to the quantum algorithms Google is using. At best it's a flashy distraction from the real work that needs to take place. At worst it's misleading, because it could make people think quantum computers can beat classical ones at anything rather than at one very narrow task. "'Supremacy' is an English word that it's going to be impossible for the public not to misinterpret," says Gil.

Google, of course, sees it rather differently.

Enter the upstart

Google was a precocious eight-year-old company when it first began tinkering with quantum problems in 2006, but it didn't form a dedicated quantum lab until 2012—the same year John Preskill, a physicist at Caltech, coined the term "quantum supremacy."

The head of the lab is Hartmut Neven, a German computer scientist with a commanding presence and a penchant for Burning Man–style chic; I saw him once in a furry blue coat and another time in an all-silver outfit that

Summit (left), with 250 petabytes of storage, is big enough to hold the complete quantum state of Sycamore (right).

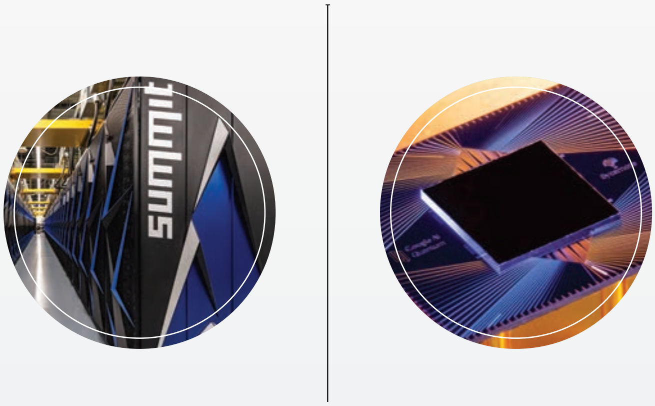

CARLOS JONES/ORNL; ERIK LUCERO/GOOGLE

A grand experiment: Quantum theory and practice

A quantum computer's basic building block is the quantum bit, or qubit. In a classical computer, a bit can store either a *0* or a *1*. A qubit can store not only *0* or *1* but also an in-between state called a superposition—which can assume lots of different values. One analogy is that if information were color, then a classical bit could be either black or white. A qubit when it's in superposition could be any color on the spectrum, and could also vary in brightness.

The upshot is that a qubit can store and process a vast quantity of information compared with a bit—and capacity increases exponentially as you connect qubits together. Storing all the information in the 53 qubits on Google's Sycamore chip would take about 72 petabytes (72 billion gigabytes) of classical computer memory. It doesn't take a lot more qubits before you'd need a classical computer the size of the planet.

But it's not straightforward. Delicate and easily disturbed, qubits need to be almost perfectly isolated from heat, vibration, and stray atoms—hence the "chandelier" refrigerators in Google's quantum lab. Even then, they can function for at most a few hundred microseconds before they "decohere" and lose their superposition.

And quantum computers aren't always faster than classical ones. They're just different, faster at some things and slower at others, and require different kinds of software. To compare their performance, you have to write a classical program that approximately simulates the quantum one.

For its experiment, Google chose a benchmarking test called "random quantum circuit sampling." It generates millions of random numbers, but with slight statistical biases that are a hallmark of the quantum algorithm. If Sycamore were a pocket calculator, it would be the equivalent of pressing buttons at random and checking that the display showed the expected results.

Google simulated parts of this on its own massive server farms as well as on Summit, the world's biggest supercomputer, at Oak Ridge National Laboratory. The researchers estimated that completing the whole job, which took Sycamore 200 seconds, would have taken Summit approximately 10,000 years. Voilà: quantum supremacy.

So what was IBM's objection? Basically, that there are different ways to get a classical computer to simulate a quantum machine—and that the software you write, the way you chop up data and store it, and the hardware you use all make a big difference in how fast the simulation can run. IBM said Google assumed the simulation would need to be cut up into a lot of chunks, but Summit, with 280 petabytes of storage, is big enough to hold the complete state of Sycamore at once. (And IBM built Summit, so it should know.)

IBM

How to program a quantum computer

At its most basic level, the software in classical computers is a sequence of logic gates like NOT, OR, and NAND that change the contents (*0* or *1*) of bits. Quantum software, similarly, consists of sequences of logic gates acting on qubits, but it has a larger and more exotic set of gates with names like SWAP (which swaps the values of two qubits around), Pauli-X (a quantum version of the NOT gate, which flips a qubit's value), and Hadamard (which turns a qubit from either *0* or *1* into a superposition of *0* and *1*). There are as yet no quantum equivalents of higher-level languages like C++ or Java, but both Google and IBM have created graphical interfaces, like the one pictured above, to make programming with gates easy.

made him look like a grungy astronaut. ("My wife buys these things for me," he explained.) Initially, Neven bought a machine built by an outside firm, D-Wave, and spent a while trying to achieve quantum supremacy on it, but without success. He says he convinced Larry Page, Google's then CEO, to invest in building quantum computers in 2014 by promising him that Google would take on Preskill's challenge: "We told him, 'Listen, Larry, in three years we will come back and put a prototype chip on your table that can at least compute a problem that is beyond the abilities of classical machines.'"

Lacking IBM's quantum expertise, Google hired a team from outside, led by John Martinis, a physicist at the University of California, Santa Barbara. Martinis and his group were already among the world's best quantum computer makers—they had managed to string up to nine qubits together—and Neven's promise to Page seemed like a worthy goal for them to aim for.

The three-year deadline came and went as Martinis's team struggled to make a chip both big enough and stable enough for the challenge. In 2018 Google released its largest processor yet, Bristlecone. With 72 qubits, it was well ahead of anything its rivals had made, and Martinis predicted it would attain quantum supremacy that same year. But a few of the team members had been working in parallel on a different chip architecture, called Sycamore, that ultimately proved able to do more with fewer qubits. Hence it was a 53-qubit chip—originally 54, but one of them malfunctioned—that ultimately demonstrated supremacy last fall.

For practical purposes, the program used in that demonstration is virtually useless—it generates random numbers, which isn't something you need a quantum computer for. But it generates them in a particular way that a classical computer would find very hard to replicate, thereby establishing the proof of concept (see opposite page).

Ask IBMers what they think of this achievement, and you get pained looks. "I don't like the word [supremacy], and I don't like the implications," says Jay Gambetta, a cautiously spoken Australian who heads IBM's quantum team. The problem, he says, is that it's virtually impossible to predict whether any given quantum calculation will be hard for a classical machine, so showing it in one case doesn't help you find other cases.

To everyone I spoke with outside IBM, this refusal to treat quantum supremacy as significant verges on pigheadedness. "Anybody who will ever have a commercially relevant offering—they have to show supremacy first. I think that's just basic logic," says Neven. Even Will Oliver, a mild-mannered MIT physicist who has been one of the most even-handed observers of the spat, says, "It's a very important milestone to show a quantum computer outperforming a classical computer at some task, whatever it is."

The quantum leap

Regardless of whether you agree with Google's position or IBM's, the next goal is clear, Oliver says: to build a quantum computer that can do something useful. The hope is that such machines could one day solve problems that require unfeasible amounts of brute-force computing power now, like modeling complex molecules to help discover new drugs and materials, or optimizing city traffic flows in real time to reduce congestion, or making longer-term weather predictions. (Eventually they might be capable of cracking the cryptographic codes used today to secure communications and financial transactions, though

DECOHER—ENCE

Qubits store information the way a sieve stores water; even the most stable ones "decohere," or fall out of their fragile quantum states, within a few hundred microseconds. Even before then, errors start to pile up. That means a quantum computer can do only so many sums before it grinds to a halt. Google's larger chips decohere after 30 to 40 microseconds, enough time for them to run through a sequence of up to 40 quantum logic gates. IBM's can reach up to 500 microseconds, but they also process gates more slowly.

TO BUILD A QUANTUM COMPUTER WITH THE POWER OF 1,000 QUBITS, YOU'D NEED A MILLION ACTUAL ONES.

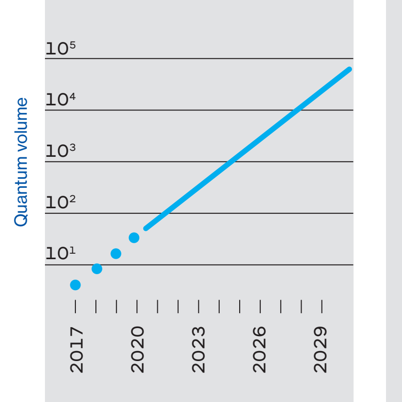

by then most of the world will probably have adopted quantum-resistant cryptography.) The trouble is that it's nearly impossible to predict what the first useful task will be, or how big a computer will be needed to perform it.

That uncertainty has to do with both hardware and software. On the hardware side, Google reckons its current chip designs can get it to somewhere between 100 and 1,000 qubits. However, just as a car's performance doesn't depend only on the size of the engine, a quantum computer's performance isn't simply determined by its number of qubits. There is a raft of other factors to take into account, including how long they can be kept from decohering, how error-prone they are, how fast they operate, and how they're interconnected. This means any quantum computer operating today reaches only a fraction of its full potential.

Software for quantum computers, meanwhile, is as much in its infancy as the machines themselves. In classical computing, programming languages are now several levels removed from the raw "machine code" that early software developers had to use, because the nitty-gritty of how data get stored, processed, and shunted around is already standardized. "On a classical computer, when you program it, you don't have to know how a transistor works," says Dave Bacon, who leads the Google team's software effort. Quantum code, on the other hand, has to be highly tailored to the qubits it will run on, so as to wring the most out of their temperamental performance. That means the code for IBM's chips won't run on those of other companies, and even techniques for optimizing Google's 53-qubit Sycamore won't necessarily do well on its future 100-qubit sibling. More important, it means nobody can predict just how hard a problem those 100 qubits will be capable of tackling.

The most anyone dares to hope for is that computers with a few hundred qubits will be cajoled into simulating some moderately complex chemistry within the next few years—perhaps even enough to advance the search for a new drug or a more efficient battery. Yet decoherence and errors will bring all these machines to a stop before they can do anything really hard like breaking cryptography.

That will require a "fault-tolerant" quantum computer, one that can compensate for errors and keep itself running indefinitely, just as classical ones do. The expected solution will be to create redundancy: make hundreds of qubits act as one, in a shared quantum state. Collectively, they can correct for individual qubits' errors. And as each qubit succumbs to decoherence, its neighbors will bring it back to life, in a never-ending cycle of mutual resuscitation.

The typical prediction is that it would take as many as 1,000 conjoined qubits to attain that stability—meaning that to build a computer with the power of 1,000 qubits, you'd need a million actual ones. Google "conservatively" estimates it can build a million-qubit processor within 10 years, Neven says, though there are some big technical hurdles to overcome, including one in which IBM may yet have the edge over Google (see opposite page).

By that time, a lot may have changed. The superconducting qubits Google and IBM currently use might prove to be the vacuum tubes of their era, replaced by something much more stable and reliable. Researchers around the world are experimenting with various methods of making qubits, though few are advanced enough to build working computers with. Rival

A new Moore's Law?

Rather than counting qubits, IBM tracks what it calls "quantum volume," a measure of how much complexity a computer can actually handle. Its goal is to keep this measure doubling every year—a quantum version of the famous Moore's Law that IBM has dubbed "Gambetta's Law," after Jay Gambetta, its chief quantum theoretician. So far, it's held for three years. That's as much data as Gordon Moore had when he postulated Moore's Law in 1965.

SOURCE: IBM

startups such as Rigetti, IonQ, or Quantum Circuits might develop an edge in a particular technique and leapfrog the bigger companies.

But given their size and wealth, both Google and IBM have a shot at becoming serious players in the quantum computing business. Companies will rent their machines to tackle problems the way they currently rent cloud-based data storage and processing power from Amazon, Google, IBM, or Microsoft. And what started as a battle between physicists and computer scientists will evolve into a contest between business services divisions and marketing departments.

Which company is best placed to win that contest? IBM, with its declining revenues, may have a greater sense of urgency than Google. It knows from bitter experience the costs of being slow to enter a market: last summer, in its most expensive purchase ever, it forked over $34 billion for Red Hat, an open-source cloud services provider, in an attempt to catch up to Amazon and Microsoft in that field and reverse its financial fortunes. Its strategy of putting its quantum machines on the cloud and building a paying business from the get-go seems designed to give it a head start.

Google recently began to follow IBM's example, and its commercial clients now include the US Department of Energy, Volkswagen, and Daimler. The reason it didn't do this sooner, says Martinis, is simple: "We didn't have the resources to put it on the cloud." But that's another way of saying it had the luxury of not having to make business development a priority.

Whether that decision gives IBM an edge is too early to say, but probably more important will be how the two companies apply their other strengths to the problem in the coming years. IBM, says Gil, will benefit from its "full stack" expertise in everything from materials science and chip fabrication to serving big corporate clients. Google, on the other hand, can boast a Silicon Valley–style culture of innovation and plenty of practice at rapidly scaling up operations.

As for quantum supremacy itself, it will be an important moment in history, but that doesn't mean it will be a decisive one. After all, everyone knows about the Wright brothers' first flight, but can anybody remember what they did afterwards? ◼

Gideon Lichfield is the editor in chief of MIT Technology Review.

A tale of two transmons

Google's and IBM's transmon qubits are almost identical, with one small but potentially crucial difference.

In both Google's and IBM's quantum computers, the qubits themselves are controlled by microwave pulses. Tiny fabrication defects mean that no two qubits respond to pulses of exactly the same frequency. There are two solutions to this: vary the frequency of the pulses to find each qubit's sweet spot, like jiggling a badly cut key in a lock until it opens; or use magnetic fields to "tune" each qubit to the right frequency.

IBM uses the first method; Google uses the second. Each approach has pluses and minuses. Google's tunable qubits work faster and more precisely, but they're less stable and require more circuitry. IBM's fixed-frequency qubits are more stable and simpler, but run more slowly.

From a technical point of view, it's pretty much a toss-up, at least at this stage. In terms of corporate philosophy, though, it's the difference between Google and IBM in a nutshell—or rather, in a qubit.

Google chose to be nimble. "In general our philosophy goes a little bit more to higher controllability at the expense of the numbers that people typically look for," says Hartmut Neven.

IBM, on the other hand, chose reliability. "There's a huge difference between doing a laboratory experiment and publishing a paper, and putting a system up with, like, 98% reliability where you can run it all the time," says Dario Gil.

Right now, Google has the edge. As machines get bigger, though, the advantage may flip to IBM. Each qubit is controlled by its own individual wires; a tunable qubit requires one extra wire. Figuring out the wiring for thousands or millions of qubits will be one of the toughest technical challenges the two companies face; IBM says it's one of the reasons they went with the fixed-frequency qubit. Martinis, the head of the Google team, says he's personally spent the past three years trying to find wiring solutions. "It's such an important problem that *I* worked on it," he jokes.

The Kuzu family at home in Cambridge, Massachusetts.

IF DNA IS LIKE SOFTWARE, CAN WE JUST FIX THE CODE? IN A RACE TO CURE HIS DAUGHTER'S ULTRA-RARE DISEASE, A GOOGLE PROGRAMMER ENTERS THE WORLD OF INDIVIDUALIZED DRUGS.

BY Erika Check Hayden
PHOTOGRAPHS BY Matthew Monteith

When you first meet her, you won't be able to tell that Ipek Kuzu suffers from a rare genetic disease. The three-year-old plays happily on her own for hours, driving her toy cars and "cooking" in her pretend kitchen. But she's not well. She's a little wobbly on her feet and doesn't say much, and if nothing is done, she may die by her mid-20s. Ipek has ataxia-telangiectasia, or A-T, a disease caused by an error in her DNA. It causes the loss of brain cells, along with a high risk of infection and cancer.

It's the sort of problem that makes doctors shake their heads. But Ipek's father, Mehmet, and mother, Tugba, hope she'll escape that fate. Thanks in part to the persistence of Mehmet, a programmer at Google, in January she became one of the first handful of US patients to receive a hyper-personalized gene medicine, tailored to treat a unique mutation. The one-person drug, designed for her by a Boston doctor, Timothy Yu, is being called "atipeksen," for "A-T" and "Ipek."

To create atipeksen, Yu borrowed from recent biotech successes like gene therapy. Some new drugs, including cancer therapies, treat disease by directly manipulating genetic information inside a patient's cells. Now doctors like Yu find they can alter those treatments as if they were digital programs. Change the code, reprogram the drug, and there's a chance of treating many genetic diseases, even those as unusual as Ipek's.

The new strategy could in theory help millions of people living with rare diseases, the vast majority of which are caused by genetic typos and have no treatment. US regulators say last year they fielded more than 80 requests to allow genetic treatments for individuals or very small groups, and that they may take steps to make tailor-made medicines easier to try. New technologies, including custom gene-editing treatments using CRISPR, are coming next.

"I never thought we would be in a position to even contemplate trying to help these patients," says Stanley Crooke, a biotechnology entrepreneur and founder of Ionis Pharmaceuticals, based in Carlsbad, California. "It's an astonishing moment."

ANTISENSE DRUG

Right now, though, insurance companies won't pay for individualized gene drugs, and no company is making them (though some plan to). Only a few patients have ever gotten them, usually after heroic feats of arm-twisting and fundraising. And it's no mistake that programmers like Mehmet Kuzu, who works on data privacy, are among the first to pursue individualized drugs. "As computer scientists, they get it. This is all code," says Ethan Perlstein, chief scientific officer at the Christopher and Dana Reeve Foundation.

A nonprofit, the A-T Children's Project, funded most of the cost of designing and making Ipek's drug. For Brad Margus, who created the foundation in 1993 after his two sons were diagnosed with A-T, the change between then and now couldn't be more dramatic. "We've raised so much money, we've funded so much research, but it's so frustrating that the biology just kept getting more and more complex," he says. "Now, we're suddenly presented with this opportunity to just fix the problem at its source."

Ipek was only a few months old when her father began looking for a cure. A geneticist friend sent him a paper describing a possible treatment for her exact form of A-T, and Kuzu flew from Sunnyvale, California, to Los Angeles to meet the scientists behind the research. But they said no one had tried the drug in people: "We need many more years to make this happen," they told him.

Kuzu didn't have years. After he returned from Los Angeles, Margus handed him a thumb drive with a video of a talk by Yu, a doctor at Boston Children's Hospital, who described how he planned to treat a young girl with Batten disease (a different neurodegenerative condition) in what press reports would later dub "a stunning illustration of personalized genomic medicine." Kuzu realized Yu was using the very same gene technology the Los Angeles scientists had dismissed as a pipe dream.

That technology is called "antisense." Inside a cell, DNA encodes information to make proteins. Between the DNA and the protein, though, come messenger molecules called RNA that ferry the gene information out of the nucleus. Think of antisense as mirror-image molecules that stick to specific RNA messages, letter for letter, blocking them from being made into proteins. It's possible to silence a gene this way, and sometimes to overcome errors, too.

Though the first antisense drugs appeared 20 years ago, the concept achieved its first blockbuster success only in 2016. That's when a drug called nusinersen, made by Ionis, was approved to treat children with spinal muscular atrophy, a genetic disease that would otherwise kill them by their second birthday.

Yu, a specialist in gene sequencing, had not worked with antisense before, but once he'd identified the genetic error causing Batten disease in his young patient, Mila Makovec, it became apparent to him he didn't have to stop there. If he knew the gene error, why not create a gene drug? "All of a sudden a lightbulb went off," Yu says. "Couldn't one try to reverse this? It was such an appealing idea, and such a simple idea, that we basically just found ourselves unable to let that go."

Yu admits it was bold to suggest his idea to Mila's mother, Julia Vitarello. But he was not starting from scratch. In a demonstration of how modular biotech drugs may become, he based milasen on the same chemistry backbone as the Ionis drug, except he made Mila's particular mutation the genetic target. Where it had taken decades for Ionis to perfect a drug, Yu now set a record: it took only eight months for him to make milasen, try it on animals, and convince the US Food and Drug Administration to let him inject it into Mila's spine.

Ipek, right, may not survive past her 20s without treatment.

Timothy Yu, below, of Boston Children's Hospital.

COURTESY PHOTO (YU)

Where it had taken decades for Ionis to perfect its drug, Yu now set a record: it took only eight months for Yu to make milasen, try it on animals, and convince the US Food and Drug Administration to let him inject it into Mila's spine.

"What's different now is that someone like Tim Yu can develop a drug with no prior familiarity with this technology," says Art Krieg, chief scientific officer at Checkmate Pharmaceuticals, based in Cambridge, Massachusetts.

SOURCE CODE

As word got out about milasen, Yu heard from more than a hundred families asking for his help. That's put the Boston doctor in a tough position. Yu has plans to try antisense to treat a dozen kids with different diseases, but he knows it's not the right approach for everyone, and he's still learning which diseases might be most amenable. And nothing is ever simple—or cheap. Each new version of a drug can behave differently and requires costly safety tests in animals.

Kuzu had the advantage that the Los Angeles researchers had already shown antisense might work. What's more, Margus agreed that the A-T Children's Project would help fund the research. But it wouldn't be fair to make the treatment just for Ipek if the foundation was paying for it. So Margus and Yu decided to test antisense drugs in the cells of three young A-T patients, including Ipek. Whichever kid's cells responded best would get picked.

While he waited for the test results, Kuzu raised about $200,000 from friends and coworkers at Google. One day, an email landed in his in-box from another Google employee who was fundraising to help a sick child. As he read it, Kuzu felt a jolt of recognition: his coworker, Jennifer Seth, was also working with Yu.

Seth's daughter Lydia was born in December 2018. The baby, with beautiful chubby cheeks, carries a mutation that causes seizures and may lead to severe disabilities. Seth's husband Rohan, a well-connected Silicon Valley entrepreneur, refers to the problem as a "tiny random mutation" in her "source code." The Seths have raised more than $2 million. Among their biggest donors: Google cofounder Sergey Brin.

CUSTOM DRUG

By then, Yu was ready to give Kuzu the good news: Ipek's cells had responded the best. So last September the family packed up and moved from California to Cambridge, Massachusetts, so Ipek could start getting atipeksen. The toddler got her first dose this January, under general anesthesia, through a lumbar puncture into her spine.

After a year, the Kuzus hope to learn whether or not the drug is helping. Doctors are measuring levels of a protein called neurofilament in Ipek's cerebrospinal fluid as a readout of how her disease is progressing. And Kuzu says a team at Johns Hopkins will compare her movements with those of other kids, both with and without A-T, to observe whether the expected disease symptoms are delayed.

One serious challenge facing gene drugs for individuals is that short of a healing miracle, it may ultimately be impossible to be sure they really work. That's because the speed with which diseases like A-T progress can vary widely from person to person. Proving a drug is effective, or revealing that it's a dud, almost always requires collecting data from many patients, not just one. "It's important for parents who are ready to pay anything, try anything, to appreciate that experimental treatments often don't work," says Holly Fernandez Lynch, a lawyer and ethicist at the University of Pennsylvania. "There are risks. Trying one could foreclose other options and even hasten death."

Kuzu says his family weighed the risks and benefits. "Since this is the first time for this kind of drug, we were a little scared," he says. But, he concluded, "there's nothing else to do. This is the only thing that might give hope to us and the other families."

Another obstacle to ultra-personal drugs is that insurance won't pay for them. And so far, pharmaceutical companies aren't interested either. They prioritize drugs that can be sold thousands of times, but as far as anyone knows, Ipek is the only person alive with her exact mutation. That leaves families facing extraordinary financial demands that only the wealthy, lucky, or well connected can meet. Developing Ipek's treatment has already cost $1.9 million, Margus estimates.

Some scientists think agencies such as the US National Institutes of Health should help fund the research, and will press their case at a meeting in Bethesda, Maryland, in April. Help could also come from the Food and Drug Administration, which is developing guidelines that may speed the work of doctors like Yu. The agency will receive updates on Mila and other patients if any of them experience severe side effects.

The FDA is also considering giving doctors more leeway to modify genetic drugs to try in new patients without securing new permissions each time. Peter Marks, director of the FDA's Center for Biologics Evaluation and Research, likens traditional drug manufacturing to factories that mass-produce identical T-shirts. But, he points out, it's now possible to order an individual basic T-shirt embroidered with a company logo. So drug manufacturing could become more customized too, Marks believes.

Custom drugs carrying exactly the message a sick kid's body needs? If we get there, credit will go to companies like Ionis that developed the new types of gene medicine. But it should also go to the Kuzus—and to Brad Margus, Rohan Seth, Julia Vitarello, and all the other parents who are trying save their kids. In doing so, they are turning hyper-personalized medicine into reality. T

Erika Check Hayden is director of the science communication program at the University of California, Santa Cruz.

MIT Technology
9:14
MIT Technology Review
Blockchain
Why it's time to start talking about blockchain ethics
Blockchain technology is changing the nature of money and organizations. We should probably start pondering the potential consequences.
01. Facebook's digital currency may force central banks to create their own October 2019
02. Facebook's Libra: Three things we don't know about the digital currency June 2019
Space Exploration Oct 11
Virgin Orbit says it wants to send spacecraft to Mars in 2022
Read more

2030

PREDICTING THE NEXT BREAKTHROUGHS

The future technologies that we imagine and choose to create are shaped by our predictions. How do we see the world in the decade to come? What problems will be most urgent? What technologies will we want? Which ones will be practical? It's why getting predictions right matters.

FIVE FORCES THAT WILL SHAPE THE FUTURE

By **Tate Ryan-Mosley**

To predict which technologies will be successful, you need to understand how our lives are changing. These are the big trends of the coming decades.

DATA EXPLOSION

We're going to need better storage, processing, and privacy.

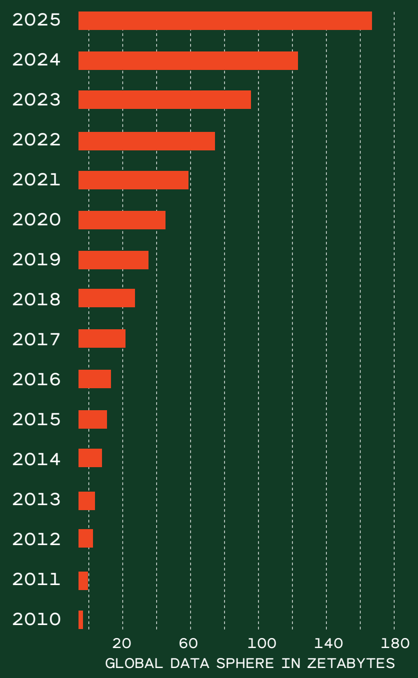

SOURCE: IDC RESEARCH, THE DIGITIZATION OF THE WORLD – FROM EDGE TO CORE (2018)

RISE IN AVERAGE GLOBAL SURFACE TEMPERATURES

As surface temperatures increase, so will sea levels, extreme storms, and habitat disruption.

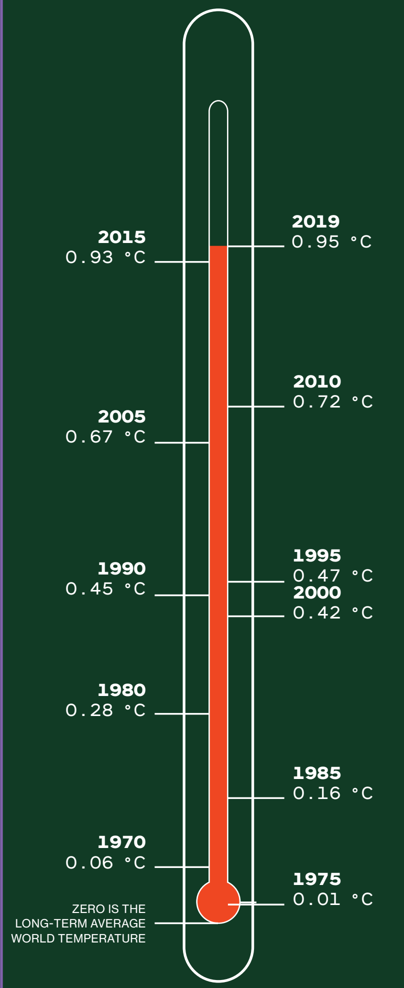

SOURCE: NOAA NATIONAL CENTERS FOR ENVIRONMENTAL INFORMATION, CLIMATE AT A GLANCE (2020)

US WEALTH GAP

Since 2007 the bottom 50% has had zero or negative wealth (i.e., debt).

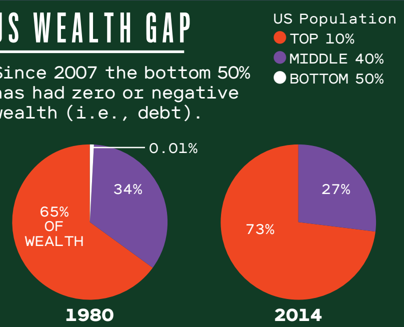

SOURCE: WORLD INEQUALITY DATABASE (2018)

AN OLDER POPULATION

Today, 9% of the global population is over 65. That's going to grow in the next decades, redefining work, health care, and our economy.

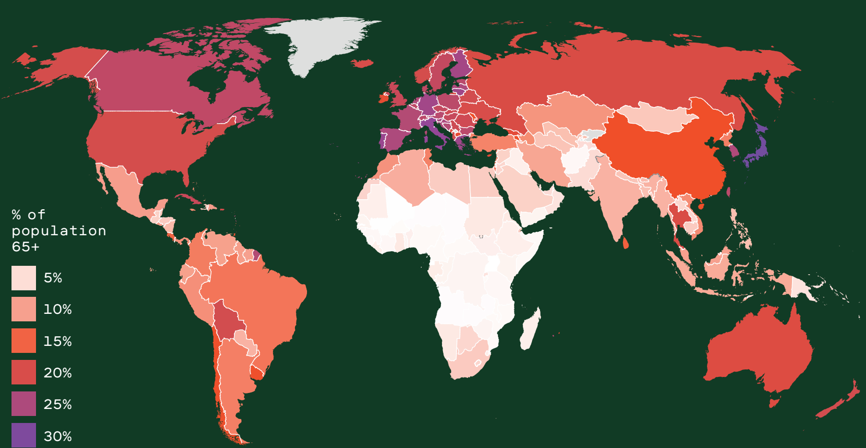

SOURCE: UNITED NATIONS, DEPARTMENT OF ECONOMIC AND SOCIAL AFFAIRS, POPULATION DIVISION (2019).

LANGUAGE EXTINCTION

FROM 1950 TO 2010, 230 LANGUAGES WENT EXTINCT. TODAY, A THIRD OF THE WORLD'S LANGUAGES HAVE FEWER THAN 1,000 SPEAKERS LEFT.

SOURCE: UNESCO WORLD LANGUAGE ATLAS (2010); ETHNOLOGUE: LANGUAGES OF THE WORLD (2019)

THE UNPREDICTABLES

Everywhere from business to medicine to the climate, forecasting the future is a complex and absolutely critical job. So how do you do it—and what comes next?

INEZ FUNG // *Professor of atmospheric science, University of California, Berkeley*

I've spoken to people who want climate model information, but they're not really sure what they're asking me for. So I say to them, "Suppose I tell you that some event will happen with a probability of 60% in 2030. Will that be good enough for you, or will you need 70%? Or would you need 90%? What level of information do you want out of climate model projections in order to be useful?"

I joined Jim Hansen's group in 1979, and I was there for all the early climate projections. And the way we thought about it then, those things are all still totally there. What we've done since then is add richness and higher resolution, but the projections are really grounded in the same kind of data, physics, and observations.

Still, there are things we're missing. We still don't have a real theory of precipitation, for example. But there are two exciting things happening there. One is the availability of satellite observations: looking at the cloud is still not totally utilized. The other is that there used to be no way to get regional precipitation patterns through history—and now there is. Scientists found these caves in China and elsewhere, and they go in, look for a nice little chamber with stalagmites, and then they chop them up and send them back to the lab, where they do fantastic uranium-thorium dating and measure oxygen isotopes in calcium carbonate. From there they can interpret a record of historic rainfall. The data are incredible: we have got over half a million years of precipitation records all over Asia.

I don't see us reducing fossil fuels by 2030. I don't see us reducing CO_2 or atmospheric methane. Some 1.2 billion people in the world right now have no access to electricity, so I'm looking forward to the growth in alternative energy going to parts of the world that have no electricity. That's important because it's education, health, everything associated with a Western standard of living. That's where I'm putting my hopes.

LEAH FASTEN

PREDICTION FOR 2030: WE'LL LIGHT UP THE WORLD ... SAFELY

Prediction for 2030: Adults will learn to grasp new ideas

ANNE LISE KJAER // *Futurist, Kjaer Global, London*

As a kid I wanted to become an archaeologist, and I did in a way. Archaeologists find artifacts from the past and try to connect the dots and tell a story about how the past might have been. We do the same thing as futurists; we use artifacts from the present and try to connect the dots into interesting narratives in the future.

When it comes to the future, you have two choices. You can sit back and think "It's not happening to me" and build a great big wall to keep out all the bad news. Or you can build windmills and harness the winds of change.

A lot of companies come to us and think they want to hear about the future, but really it's just an exercise for them – let's just tick that box, do a report, and put it on our bookshelf.

So we have a little test for them. We do interviews, we ask them questions; then we use a model called a Trend Atlas that considers both the scientific dimensions of society and the social ones. We look at the trends in politics, economics, societal drivers, technology, environment, legislation – how does that fit with what we know currently? We look back maybe 10, 20 years: can we see a little bit of a trend and try to put that into the future?

What's next? Obviously with technology we can educate much better than we could in the past. But it's a huge opportunity to educate the parents of the next generation, not just the children. Kids are learning about sustainability goals, but what about the people who actually rule our world?

Prediction for 2030: We'll get better at being uncertain

PHILIP TETLOCK // *Coauthor of* Superforecasting *and professor, University of Pennsylvania*

At the Good Judgment Project, we try to track the accuracy of commentators and experts in domains in which it's usually thought impossible to track accuracy. You take a big debate and break it down into a series of testable short-term indicators. So you could take a debate over whether strong forms of artificial intelligence are going to cause major dislocations in white-collar labor markets by 2035, 2040, 2050. A lot of discussion already occurs at that level of abstraction—but from our point of view, it's more useful to break it down and to say: If we were on a long-term trajectory toward an outcome like that, what sorts of things would we expect to observe in the short term? So we started this off in 2015, and in 2016 AlphaGo defeated people in Go. But then other things didn't happen: driverless Ubers weren't picking people up for fares in any major American city at the end of 2017. Watson didn't defeat the world's best oncologists in a medical diagnosis tournament. So I don't think we're on a fast track toward the singularity, put it that way.

Forecasts have the potential to be either self-fulfilling or self-negating—Y2K was arguably a self-negating forecast. But it's possible to build that into a forecasting tournament by asking conditional forecasting questions: i.e., How likely is X conditional on our doing this or doing that?

What I've seen over the last 10 years, and it's a trend that I expect will continue, is an increasing openness to the quantification of uncertainty. I think there's a grudging, halting, but cumulative movement toward thinking about uncertainty, and more granular and nuanced ways that permit keeping score.

DVORA PHOTOGRAPHY (KJAER) / COURTESY PHOTO (TETLOCK)

RYAN YOUNG

PREDICTION FOR 2030: WE'LL BE MORE—AND LESS—PRIVATE

KEITH CHEN // *Associate professor of economics, UCLA*

When I worked on Uber's surge pricing algorithm, the problem it was built to solve was very coarse: we were trying to convince drivers to put in extra time when they were most needed. There were predictable times—like New Year's—when we knew we were going to need a lot of people. The deeper problem was that this was a system with basically no control. It's like trying to predict the weather. Yes, the amount of weather data that we collect today—temperature, wind speed, barometric pressure, humidity data—is 10,000 times greater than what we were collecting 20 years ago. But we still can't predict the weather 10,000 times further out than we could back then. And social movements—even in a very specific setting, such as where riders want to go at any given point in time—are, if anything, even more chaotic than weather systems.

These days what I'm doing is a little bit more like forensic economics. We look to see what we can find and predict from people's movement patterns. We're just using simple cell-phone data like geolocation, but even just from movement patterns, we can infer salient information and build a psychological dimension of you. What terrifies me is I feel like I have much worse data than Facebook does. So what are they able to understand with their much better information?

I think the next big social tipping point is people actually starting to really care about their privacy. It'll be like smoking in a restaurant: it will quickly go from causing outrage when people want to stop it to suddenly causing outrage if somebody does it. But at the same time, by 2030 almost every Chinese citizen will be completely genotyped. I don't quite know how to reconcile the two.

PREDICTION FOR 2030: WE'RE GOING TO SEE A LOT MORE HUMBLE TECHNOLOGY

SARAH DERAGON

ANNALEE NEWITZ // *Science fiction and nonfiction author, San Francisco*

Every era has its own ideas about the future. Go back to the 1950s and you'll see that people fantasized about flying cars. Now we imagine bicycles and green cities where cars are limited, or where cars are autonomous. We have really different priorities now, so that works its way into our understanding of the future.

Science fiction writers can't actually make predictions. I think of science fiction as engaging with questions being raised in the present. But what we can do, even if we can't say what's definitely going to happen, is offer a range of scenarios informed by history.

There are a lot of myths about the future that people believe are going to come true right now. I think a lot of people—not just science fiction writers but people who are working on machine learning—believe that relatively soon we're going to have a human-equivalent brain running on some kind of computing substrate. This is as much a reflection of our time as it is what might actually happen.

It seems unlikely that a human-equivalent brain in a computer is right around the corner. But we live in an era where a lot of us feel like we live inside computers already, for work and everything else. So of course we have fantasies about digitizing our brains and putting our consciousness inside a machine or a robot.

I'm not saying that those things could never happen. But they seem much more closely allied to our fantasies in the present than they do to a real technical breakthrough on the horizon.

We're going to have to develop much better technologies around disaster relief and emergency response, because we'll be seeing a lot more floods, fires, storms. So I think there is going to be a lot more work on really humble technologies that allow you to take your community off the grid, or purify your own water. And I don't mean in a creepy survivalist way; I mean just in a this-is-how-we-are-living-now kind of way.

Prediction for 2030: Humans and machines will make decisions together

FINALE DOSHI-VELEZ // *Associate professor of computer science, Harvard*

In my lab, we're trying to answer questions like "How might this patient respond to this antidepressant?" or "How might this patient respond to this vasopressor?" So we get as much data as we can from the hospital. For a psychiatric patient, we might have everything about their heart disease, kidney disease, cancer; for a blood pressure management recommendation for the ICU, we have all their oxygen information, their lactate, and more.

Some of it might be relevant to making predictions about their illnesses, some not, and we don't know which is which. That's why we ask for the large data set with everything.

There's been about a decade of work trying to get unsupervised machine-learning models to do a better job at making these predictions, and none worked really well. The breakthrough for us was when we found that all the previous approaches for doing this were wrong in the exact same way. Once we untangled all of this, we came up with a different method.

We also realized that even if our ability to predict what drug is going to work is not always that great, we can more reliably predict what drugs are not going to work, which is almost as valuable.

I'm excited about combining humans and AI to make predictions. Let's say your AI has an error rate of 70% and your human is also only right 70% of the time. Combining the two is difficult, but if you can fuse their successes, then you should be able to do better than either system alone. How to do that is a really tough, exciting question.

All these predictive models were built and deployed and people didn't think enough about potential biases. I'm hopeful that we're going to have a future where these human-machine teams are making decisions that are better than either alone.

NOAH WILLMAN (DOSHI-VELEZ) / GUILLAUME SIMONEAU (DIALLO)

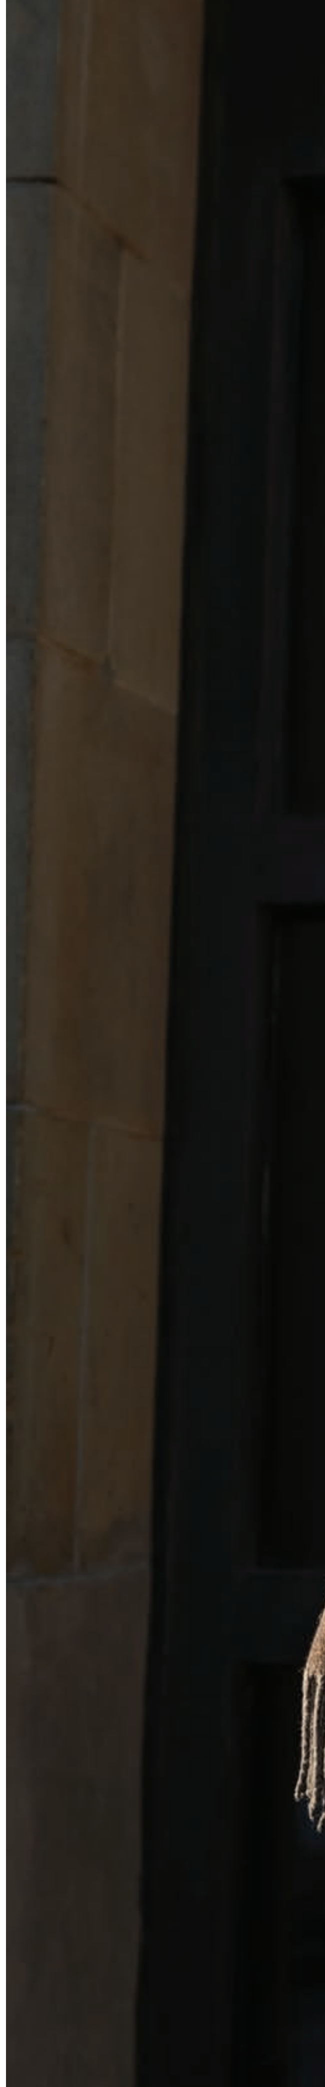

PREDICTION FOR 2030: MACHINE-BASED FORECASTING WILL BE REGULATED

ABDOULAYE BANIRE DIALLO // *Professor, director of the bioinformatics lab, University of Quebec at Montreal*

When a farmer in Quebec decides whether to inseminate a cow or not, it might depend on the expectation of milk that will be produced every day for one year, two years, maybe three years after that. Farms have management systems that capture the data and the environment of the farm. I'm involved in projects that add a layer of genetic and genomic data to help forecasting—to help decision makers like the farmer to have a full picture when they're thinking about replacing cows, improving management, resilience, and animal welfare.

With the emergence of machine learning and AI, what we're showing is that we can help tackle problems in a way that hasn't been done before. We are adapting it to the dairy sector, where we've shown that some decisions can be anticipated 18 months in advance just by forecasting based on the integration of this genomic data. I think in some areas such as plant health we have only achieved 10% or 20% of our capacity to improve certain models.

Until now AI and machine learning have been associated with domain expertise. It's not a public-wide thing. But less than 10 years from now they will need to be regulated. I think there are a lot of challenges for scientists like me to try to make those techniques more explainable, more transparent, and more auditable.

AI STILL GETS CONFUSED ABOUT HOW THE WORLD WORKS

Artificial intelligence won't be very smart if computers don't grasp cause and effect. That's something even humans have trouble with.

BY BRIAN BERGSTEIN

ILLUSTRATION BY SAIMAN CHOW

In less than a decade, computers have become extremely good at diagnosing diseases, translating languages, and transcribing speech. They can outplay humans at complicated strategy games, create photorealistic images, and suggest useful replies to your emails.

Yet despite these impressive achievements, artificial intelligence has glaring weaknesses.

Machine-learning systems can be duped or confounded by situations they haven't seen before. A self-driving car gets flummoxed by a scenario that a human driver could handle easily. An AI system laboriously trained to carry out one task (identifying cats, say) has to be taught all over again to do something else (identifying dogs). In the process, it's liable to lose some of the expertise it had in the original task. Computer scientists call this problem "catastrophic forgetting."

These shortcomings have something in common: they exist because AI systems don't understand causation. They see that some events are associated with other events, but they don't ascertain which things directly make other things happen. It's as if you knew that the presence of clouds made rain likelier, but you didn't know clouds caused rain.

Understanding cause and effect is a big aspect of what we call common sense, and it's an area in which AI systems today "are clueless," says Elias Bareinboim. He should know: as the director of the new Causal Artificial Intelligence Lab at Columbia University, he's at the forefront of efforts to fix this problem.

His idea is to infuse artificial-intelligence research with insights from the relatively new science of causality, a field shaped to a huge extent by Judea Pearl, a Turing Award–winning scholar who considers Bareinboim his protégé.

As Bareinboim and Pearl describe it, AI's ability to spot correlations—e.g., that clouds make rain more likely—is merely the simplest level of causal reasoning. It's good enough to have driven the boom in the AI technique known as deep learning over the past decade. Given a great deal of data about familiar situations, this method can lead to very good predictions. A computer can calculate the probability that a patient with certain symptoms has a certain disease, because it has learned just how often thousands or even millions of other people with the same symptoms had that disease.

But there's a growing consensus that progress in AI will stall if computers don't get better at wrestling with causation. If machines could grasp that certain things lead to other things, they wouldn't have to learn everything anew all the time—they could take what they had learned in one domain and apply it to another. And if machines could use common sense we'd be able to put more trust in them to take actions on their own, knowing that they aren't likely to make dumb errors.

Today's AI has only a limited ability to infer what will result from a given action. In reinforcement learning, a technique that has allowed machines to master games like chess and Go, a system uses extensive trial and error to discern which moves will essentially cause them to win. But this approach

doesn't work in messier settings in the real world. It doesn't even leave a machine with a general understanding of how it might play other games.

An even higher level of causal thinking would be the ability to reason about why things happened and ask "what if" questions. A patient dies while in a clinical trial; was it the fault of the experimental medicine or something else? School test scores are falling; what policy changes would most improve them? This kind of reasoning is far beyond the current capability of artificial intelligence.

PERFORMING MIRACLES

The dream of endowing computers with causal reasoning drew Bareinboim from Brazil to the United States in 2008, after he completed a master's in computer science at the Federal University of Rio de Janeiro. He jumped at an opportunity to study under Judea Pearl, a computer scientist and statistician at UCLA. Pearl, 83, is a giant—*the* giant—of causal inference, and his career helps illustrate why it's hard to create AI that understands causality.

Even well-trained scientists are apt to misinterpret correlations as signs of causation—or to err in the opposite direction, hesitating to call out causation even when it's justified. In the 1950s, for example, a few prominent statisticians muddied the waters around whether tobacco caused cancer. They argued that without an experiment randomly assigning people to be smokers or nonsmokers, no one could rule out the possibility that some unknown—stress, perhaps, or some gene—caused people both to smoke and to get lung cancer.

Eventually, the fact that smoking causes cancer was definitively established, but it needn't have taken so long. Since then, Pearl and other statisticians have devised a mathematical approach to identifying what facts would be required to support a causal claim. Pearl's method shows that, given the prevalence of smoking and lung cancer, an independent factor causing both would be extremely unlikely.

Conversely, Pearl's formulas also help identify when correlations can't be used to determine causation. Bernhard Schölkopf, who researches causal AI techniques as a director at Germany's Max Planck Institute for Intelligent Systems, points out that you can predict a country's birth rate if you know its population of storks. That isn't because storks deliver babies or because babies attract storks, but probably because economic development leads to more babies and more storks. Pearl has helped give statisticians and computer scientists ways of attacking such problems, Schölkopf says.

Pearl's work has also led to the development of causal Bayesian networks—software that sifts through large amounts of data to detect which variables appear to have the most influence on other variables. For example, GNS Healthcare, a company in Cambridge, Massachusetts, uses these techniques to advise researchers about experiments that look promising.

In one project, GNS worked with researchers who study multiple myeloma, a kind of blood cancer. The researchers wanted to know why some patients with the disease live longer than others after getting stem-cell transplants, a common form of treatment. The software churned through data with 30,000 variables and pointed to a few that seemed especially likely to be causal. Biostatisticians and experts in the disease zeroed in on one in particular: the level of a certain protein in patients' bodies. Researchers could then run a targeted clinical trial to see whether patients with the protein did indeed benefit more from the treatment. "It's way faster than poking here and there in the lab," says GNS cofounder Iya Khalil.

Nonetheless, the improvements that Pearl and other scholars have achieved in causal theory haven't yet made many inroads in deep learning, which identifies correlations without too much worry about causation. Bareinboim is working to take the next step: making computers more useful tools for human causal explorations.

ELIAS BAREINBOIM: AI SYSTEMS ARE CLUELESS WHEN IT COMES CAUSATION.

JUDEA PEARL: HIS THEORY OF CAUSAL REASONING HAS TRANSFORMED SCIENCE.

One of his systems, which is still in beta, can help scientists determine whether they have sufficient data to answer a causal question. Richard McElreath, an anthropologist at the Max Planck Institute for Evolutionary Anthropology, is using the software to guide research into why humans go through menopause (we are the only apes that do).

The hypothesis is that the decline of fertility in older women benefited early human societies because women who put more effort into caring for grandchildren ultimately had more descendants. But what evidence might exist today to support the claim that children do better with grandparents around? Anthropologists can't just compare the educational or medical outcomes of children who have lived with grandparents and those who haven't. There are what statisticians call confounding factors: grandmothers might be likelier to live with grandchildren who need the most help. Bareinboim's software can help McElreath discern which studies about kids who grew up with their grandparents are least riddled with confounding factors and could be valuable in answering his causal query. "It's a huge step forward," McElreath says.

THE LAST MILE

Bareinboim talks fast and often gestures with two hands in the air, as if he's trying to balance two sides of a mental equation. It was halfway through the semester when I visited him at Columbia in October, but it seemed as if he had barely moved into his office—hardly anything on the walls, no books on the shelves, only a sleek Mac computer and a whiteboard so dense with equations and diagrams that it looked like a detail from a cartoon about a mad professor.

He shrugged off the provisional state of the room, saying he had been very busy giving talks about both sides of the causal revolution. Bareinboim believes work like his offers the opportunity not just to incorporate causal thinking into machines, but also to improve it in humans.

Getting people to think more carefully about causation isn't necessarily much easier than teaching it to machines, he says. Researchers in a wide range of disciplines, from molecular biology to public policy, are sometimes content to unearth correlations that are not actually rooted in causal relationships. For instance, some studies suggest drinking alcohol will kill you early, while others indicate that moderate consumption is fine and even beneficial, and still other research has found that heavy drinkers outlive nondrinkers. This phenomenon, known as the "reproducibility crisis," crops up not only in medicine and nutrition but also in psychology and economics. "You can see the fragility of all these inferences," says Bareinboim. "We're flipping results every couple of years."

He argues that anyone asking "what if"—medical researchers setting up clinical trials, social scientists developing pilot programs, even web publishers preparing A/B tests—should start not merely by gathering data but by using Pearl's causal logic and software like Bareinboim's to determine whether the available data could possibly answer a causal hypothesis. Eventually, he envisions this leading to "automated scientist" software: a human could dream up a causal question to go after, and the software would combine causal inference theory with machine-learning techniques to rule out experiments that wouldn't answer the question. That might save scientists from a huge number of costly dead ends.

Bareinboim described this vision while we were sitting in the lobby of MIT's Sloan School of Management, after a talk he gave last fall. "We have a building here at MIT with, I don't know, 200 people," he said. How do those social scientists, or any scientists anywhere, decide which experiments to pursue and which data points to gather? By following their intuition: "They are trying to see where things will lead, based on their current understanding."

That's an inherently limited approach, he said, because human scientists designing an experiment can consider only a handful of variables in their minds at once. A computer, on the other hand, can see the interplay of hundreds or thousands of variables. Encoded with "the basic principles" of Pearl's causal calculus and able to calculate what might happen with new sets of variables, an automated scientist could suggest exactly which experiments the human researchers should spend their time on. Maybe some public policy that has been shown to work only in Texas could be made to work in California if a few causally relevant factors were better appreciated. Scientists would no longer be "doing experiments in the darkness," Bareinboim said.

He also doesn't think it's that far off: "This is the last mile before the victory."

> **"Pearl says AI can't be truly intelligent until it has a rich understanding of cause and effect, which would enable the introspection that is at the core of cognition.**

WHAT IF?

Finishing that mile will probably require techniques that are just beginning to be developed. For example, Yoshua Bengio, a computer scientist at the University of Montreal who shared the 2018 Turing Award for his work on deep learning, is trying to get neural networks—the software at the heart of deep learning—to do "meta-learning" and notice the causes of things.

As things stand now, if you wanted a neural network to detect when people are dancing, you'd show it many, many images of dancers. If you wanted it to identify when people are running, you'd show it many, many images of runners. The system would learn to distinguish runners from dancers by identifying features that tend to be different in the images, such as the positions of a person's hands and arms. But Bengio points out that fundamental knowledge about the world can be gleaned by analyzing the things that are similar or "invariant" across data sets. Maybe a neural network could learn that movements of the legs physically cause both running and dancing. Maybe after seeing these examples and many others that show people only a few feet off the ground, a machine would eventually understand something about gravity and how it limits human movement. Over time, with enough meta-learning about variables that are consistent across data sets, a computer could gain causal knowledge that would be reusable in many domains.

For his part, Pearl says AI can't be truly intelligent until it has a rich understanding of cause and effect. Although causal reasoning wouldn't be sufficient for an artificial general intelligence, it's necessary, he says, because it would enable the introspection that is at the core of cognition. "What if" questions "are the building blocks of science, of moral attitudes, of free will, of consciousness," Pearl told me.

You can't draw Pearl into predicting how long it will take for computers to get powerful causal reasoning abilities. "I am not a futurist," he says. But in any case, he thinks the first move should be to develop machine-learning tools that combine data with available scientific knowledge: "We have a lot of knowledge that resides in the human skull which is not utilized."

Brian Bergstein, a former editor at MIT Technology Review, is deputy opinion editor at the Boston Globe.

GRAND

How a movement to make smart, funny, critical predictions turned into fodder for ad campaigns and TV spots.

by TIM MAUGHAN

TBD CATALOG
2014
Combines Silicon Valley fever dreams with a satiric SkyMall presentation.

MALTESE FALCON
1930
Dashiell Hammett's MacGuffin was a piece of proto-design fiction.

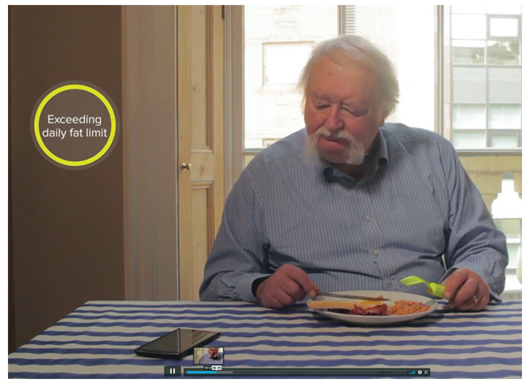

"UNINVITED GUESTS"
2015
This short film by Superflux shows an elderly man getting the better of surveillance devices.

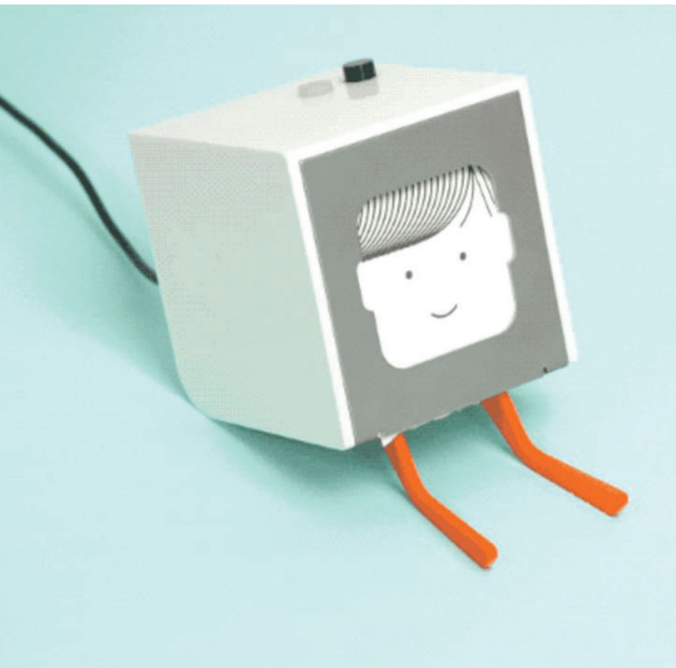

LITTLE PRINTER
2012
A design fiction idea that became a real product, Berg London's chirpy thermal printer took your feed of social media, news, and weather updates and turned it into a physical object.

DESIGNS

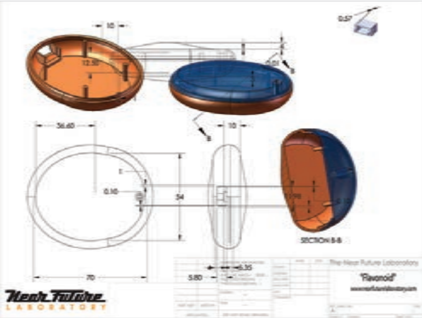

FLAVANOID
2007
A wearable device that measures your activity and uses the data to change your avatar in the virtual world Second Life.

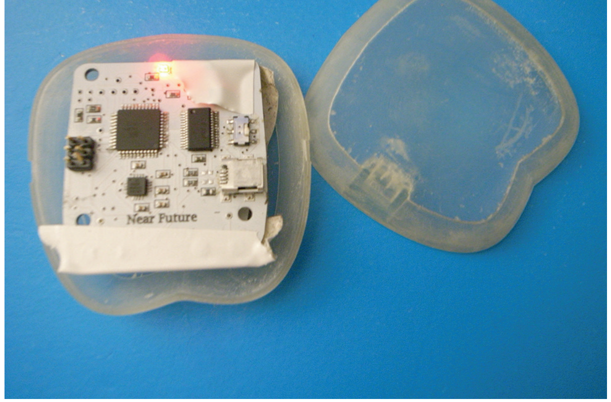

SLOW MESSENGER
2007
This gadget deliberately slows down the receipt of messages to push back against rushed, always-on culture.

JULIAN BLEECKER
"Ideas are kind of like a dime a dozen."

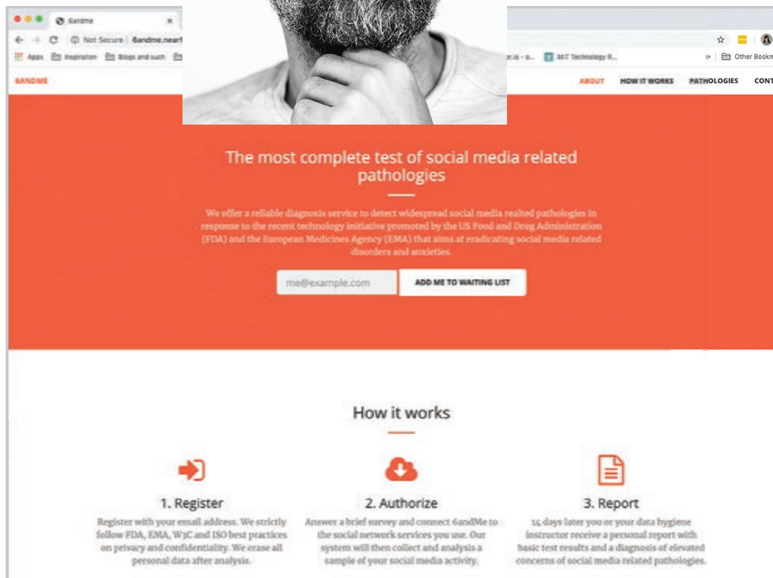

BUTTONS: BLIND CAMERA
2010
Sascha Pohflepp's digital camera has no lens: instead, it shows you a photo taken and shared by somebody else at the exact same moment.

6ANDME
2015
The service analyzes your social-media accounts to diagnose various fictional ailments.

Bruce Sterling wasn't originally meant to be part of the discussion. It was March 13, 2010, in Austin, Texas, and a small group of designers were on stage at the South by Southwest interactive festival, talking about an emerging discipline they called "design fiction."

"They asked me to join the panel at the last minute," Sterling tells me, laughing. "They knew that I'd been [involved with] South by Southwest for a long time and this would give them some cred."

A science fiction novelist who'd helped launch the cyberpunk movement in the 1980s, Sterling had actually coined the term design fiction in a 2005 book, but he hadn't exactly taken ownership of the still-nebulous concept. What happened that day made it much clearer, though, and set off an explosion of ideas for everyone in attendance.

"People went out of that room and they were kind of visibly shaken," he says. "Some guy came up in the back and told us, with this pale kind of look, 'I think I'm starting to get it.'"

The panel's organizer was Julian Bleecker, an artist, technologist, and product designer from Los Angeles. He wanted to share his work—a new practice where designers and engineers used their skills to go beyond just thinking up and prototyping new consumer products. He wanted them to create objects that were not intended to be real products but *could* have been, and use them as portals for talking about tomorrow.

"Design fiction is a mix of science fact, design and science fiction," Bleecker wrote on his blog in 2009. It "recombines the traditions of writing and storytelling with the material crafting of objects." The objects made in design fiction are "diagetic prototypes," he suggested. They are "props that help focus the imagination and speculate about possible near future worlds—whether profound change or simple, even mundane social practices."

One of the earliest examples is the late artist Sascha Pohflepp's *Buttons: Blind Camera*. Made in 2010, it is a sleek-looking digital camera that takes the minimal, post-Apple industrial design aesthetic to an extreme. It has only one button, a small color screen, and apparently no lens. Press the button and it, like any other camera, captures a moment of time in the form of a photograph. The difference is that it's not a moment of *your* time. Instead, the camera connects to the internet to find another photo taken and shared by somebody else at the exact time you pressed that button, downloads it, and displays it on the screen.

It was a brilliantly simple idea, but crucially, it was not just a piece of concept art, or a prop in a speculative movie, or an art student's mock-up. It was a real, functioning device. Pohflepp built it from the guts of a Sony Ericsson cell phone and code he'd written himself.

"It's an object that's somehow imbued with kind of a narrative function," Bleecker says. "It helps tell a story; it pushes and pulls on characters in certain ways. I think the classic example is the Maltese Falcon. Hitchcock called them MacGuffins. It's the thing around which the drama evolves and develops and moves."

In design fiction, the process of making—rather than just imagining—is the process of learning. "I don't want to dismiss the significance or importance of a good creative idea, but ideas are kind of like a dime a dozen," Bleecker says.

Back in 2007 he'd built the Slow Messenger, a handheld device that received messages but delayed presenting them—by minutes, days, or sometimes even years. It poked at the idea of instant, always-on communication that the internet was thrusting onto us. Shortly after that, he cofounded the Near Future Laboratory, a studio that produced this kind of exploratory work.

The lab created things like the TBD Catalog, a SkyMall-style magazine full of hilarious advertisements for disposable, very plausibly makeable near-future consumer crap with a tone reminiscent of Paul Verhoeven's satirical sci-fi movies *Robocop* and *Starship Troopers*. Then there is 6andMe, a service that analyzes your social-media accounts and diagnoses supposed "social media related pathologies." ("Systrom's Anxiety," named for the Instagram cofounder, is the drive to record moments of one's life for fear of not being able to repeat them in the future; "Six Degrees Jealousy" is when we envy somebody for getting more likes.) These maladies are all fictional, as is the service's analysis, but the fake reports are sinisterly familiar to anybody who has spent time nervously checking Twitter or Instagram feeds.

It's that sense of something unsettling yet just a few minutes into the future that you get from *Black Mirror* or Spike Jonze's movie *Her*.

As design fiction emerged, it turned out that governments, multinational companies, and art galleries were all interested in exploring what the future looked like, and intrigued by the charismatic objects the movement produced. The Near Future Lab joined a number of other boutique agencies that offered speculative services to their clients.

"We use objects to ask 'Why/Why Not?' questions," explains Scott Smith, one of the founders of Changeist, a consultancy now based in the Netherlands that works mainly with large institutions. "We try to use the familiar forms and language of these bureaucracies to speak back to them—manuals, maps, forms, kits, procedures, organizations, and so on."

Design fiction rapidly expanded from a practice into an aesthetic: a style that used the languages of consumer product design and advertising to create fictional objects so instantly familiar to audiences that they feel real, close, or even inevitable. It's that sense of something being unsettling yet just a few minutes into the future that you get from every dystopian app in *Black Mirror* or the ubiquitous voice assistant in Spike Jonze's movie *Her*.

As the style went mainstream and commercial, however, it started to change. In 2011, glass manufacturer Corning released "A Day Made of Glass," depicting a day in the life of a painfully perfect-looking family. Its five minutes of sleek concept video show every single glass surface—windows, mirrors, tabletops—becoming touch screens. Its 26 million YouTube views led

Marketing Daily magazine to call it “the most watched corporate video of all time.” As dazzling and high-tech as it looked on release, it feels quite dull and naïve—even dystopian—nine years later. More important, it’s utterly lacking in the anarchic, critical attitude that marked early, genuine design fiction work. It was a sign of how corporate interests would appropriate design fiction—and declaw it.

A more recent example is a May 2019 Amazon ad for the Echo smart speaker, “Caring Is Sharing.” The 30-second spot shows a young man bringing his grandfather an Echo and installing it in his apartment, presumably to keep him company and to let family members stay in touch with him. He’s grumpy about it at first, reluctant to acknowledge it, but the next time his grandson comes to visit, he’s using it happily.

Though at first glance it seems like any other TV ad, “Caring Is Sharing” looks and feels eerily similar to “Uninvited Guests,” a five-minute satirical film made by Superflux, a London-based “speculative design agency,” in 2015. That video similarly portrays an old man living on his own who has been given a range of surveillance devices by well-meaning family members: a smart fork that measures the nutrients in his food and nags him about his salt and fat intake, a smart walking cane that scolds him if he doesn’t get his recommended daily steps, and a device that connects to his bed to make sure he’s getting enough sleep. But instead of succumbing to the intrusions of these devices—as

1. A Near Future project to create a unique controller for the game Katamari Damacy.

2. Bleecker’s sketches wonder what real-world gestures are appropriate to turn into in-game actions. Could snowboarding be used to steer your character?

3. A prototype for Slow Messenger, which delays inbound mail by as much as a decade.

in the Amazon ad—the protagonist of “Uninvited Guests” finds ways to fool them. He puts the smart fork in a plate of salad while eating fish and chips, pays a local teenager in beer to walk the smart cane for him, and piles books on his bed so it looks as if he’s sleeping when he watches TV.

Superflux’s cofounder Anab Jain hadn’t seen the Amazon film when I spoke to her, but she’s aware that corporations have used the speculative approach for marketing. “It’s deeply problematic,” she says. “It’s why we say no to work more than we say yes.” Jain, who prefers the term “speculative design” or “critical design” (because “frankly, all design is fiction until it’s real”), says some prospective clients pay lip service “to the criticality and to the questioning,” but “in the end they just want a PR exercise.”

For Bleecker, this isn’t what design fiction should be. “There’s a number of those kinds of films that are essentially marketing exercises,” he says. “There was no sense that they were meant to be used internally to reflect upon and consider directions in which the company is going. They definitely come across as advertisements: ‘Look, we’re futuristic, we’ve got lots of concepts that relate to flat screens and graphics circulating and swirling around.’”

In many ways design fiction’s path from a smart, anarchic movement to a marketing language for the industries it set out to lampoon is painfully familiar.

Last year designer and artist Tobias Revell claimed that “speculative design has failed to achieve the meaningful tools for change that we once hoped for.” It had become, he said, “a whitewashing exercise” for tech companies.

Others, meanwhile, suggest it was never going to be able to achieve its original goals: it was too wrapped up in corporate hegemony from the beginning, too exclusive and elitist. Design fiction was focused on “projects that clearly reflect the fear of losing first-world privilege in bleak, dystopic futures,” wrote Brazilian design duo A Parede in 2014.

Perhaps more practically, those working in the field faced another, also familiar issue: they had to balance their desire to do critical work with their need to pay the bills. This inevitably watered down their ability to achieve distance from the organizations that were lifting their ideas and aesthetics.

For agencies like Superflux and Changeist, that means continuing to take corporate contracts and using the money to work on more personal projects. Others have taken jobs with governments or big tech themselves. But while the surface may have been captured by Hollywood and the advertising industry, some folks are still plugging away, trying to navigate a path between the critical and the corporate.

And then there’s Bleecker himself. Ten years on, he’s still running Near Future Lab, working with clients, building objects from the future, and throwing out his own brand of wild ideas. But he’s also working on Omata, a small two-person company that makes high-tech cycling accessories. Its flagship product is a $550 screenless cycling computer that looks like a giant Swiss watch. It is a product for privileged first-worlders, not a tool for change; it is a beautiful object, obviously lovingly designed and born out of Bleecker’s very personal obsessions. But it is also a deliberate challenge to the idea of what would be expected from such a device.

“It almost seemed to me like … it would have to be something unexpected,” he says.

By doing the opposite of everything that corporate technology companies might try—the antithesis of a suite of interchangeable, low-cost, shrunken-down touch-screen gizmos—Omata is rooted in design fiction, with its mission to challenge us to imagine other futures and see the world differently.

Tim Maughan is a journalist and author. His debut novel Infinite Detail was picked by The Guardian as its best science fiction book of 2019.

THE WORLD IN 2030... BY

At the World Economic Forum in Davos, the elite of the elite gather to hatch plans for the future of the planet. I asked some of this year's participants to tell me one thing they think will happen by 2030 that most people don't realize.—*Gideon Lichfield*

AI WILL CAUSE A PRODUCTIVITY BOOM

Erik Brynjolfsson, director, MIT Initiative on the Digital Economy (USA)

Machine learning has advanced tremendously over the past decade, yet US productivity growth has fallen by 50% since 2004. It's not uncommon with powerful new general-purpose technologies to see first a dip in productivity growth followed by an increase. It takes time. With the steam engine, we saw the rise of industrialization. With electricity, factories were reinvented. Computers obviously changed many aspects of society, but e-commerce is still a minority of total retail trade, 25 years after Amazon was started. Likewise, machine learning is going to take a while to propagate through the economy. What's needed is investments in new skills, and businesses that are willing to fundamentally rethink their supply chains, their relationships with customers, and the kinds of products and services they deliver. As they do that, the productivity is going to come online.

"E-commerce is still a minority of total retail trade, 25 years after Amazon was started."

AFRICA WILL BE A TEST BED FOR HUMAN-ROBOT COEXISTENCE

Wanuri Kahiu, science fiction writer and filmmaker (Kenya)

Just as Kenya has been a place where digital payment technologies took off, I think it will become a testing ground for how people interact with AI and robots. The barriers to entry are low and there are few laws or social mores around AI, so it's like a blank slate for experiments in coexistence between humans and machines. In Kinshasa almost 10 years ago, they installed robotic traffic cops and people obeyed them more than the human police, because the robots were not corrupt. There's lots of potential for localized AI applications that help Africa deal with African problems, which is important because by 2050, one in four people will be African.

CONSUMERS WILL HAVE MORE POWER AND MORE PROTECTION

Helena Leurent, director-general, Consumers International (UK)

Consumers will be part of data trusts and cooperatives that can safeguard their rights, negotiate for them on how their data is used, alert them to how they are being watched, and audit organizations that use their data. As an example, consumers might want their respective data trusts to connect directly to farmers who guarantee to use sustainable growing practices. The consumers would get better prices and have more information about what they're buying; the farmers could get data and guarantees about purchasing patterns and would be able to differentiate their products. This "agricultural data commons" could spark innovation in products and services that both give consumers more choice and lead to greater sustainability.

THE DOLLAR WILL NO LONGER BE THE WORLD'S RESERVE CURRENCY

Michael Casey, chief content officer, CoinDesk (USA)

The dollar is the reserve currency because of its stability. If companies in two different countries sign a contract with payment due in 90 days, they set the transaction in dollars to protect against exchange-rate fluctuations. But when there are digital currencies with programmable smart contracts that can convert at an agreed rate and keep the payment in escrow until it's due, they won't need the dollar any more. This means the advantages to traditional US companies will diminish, but innovative, decentralized, globally minded companies will succeed.

BUSINESS WIRE (BRYNJOLFSSON); WORLD ECONOMIC FORUM/BORIS BALDINGER (KAHIU); WORLD ECONOMIC FORUM/FARUK PINJO (LEURENT); COURTESY PHOTO (CASEY)

THE PEOPLE SHAPING IT

WE'LL RECOGNIZE THE BRITTLENESS OF 20TH-CENTURY INFRASTRUCTURE

Genevieve Bell, director, 3A Institute and senior fellow, Intel (Australia)

Over the last six weeks my country has been on fire, and I think 2030 looks like the world I'm now living in. One, the climate is changing faster and faster. Two, Australians are suddenly having to think much harder about how both their own personal data and government data is made accessible so they can get timely fire projections, evacuation requests, air-quality reports, and so on—so the questions about data that only those of us at the forefront of technology were asking are now mainstream. And three, we'll have to contend with the fact that all the infrastructures of the 20th century—electricity, water, communications, civil society itself—are brittle, and this brittleness will make the 21st century harder to deliver.

WE'LL GROW PLASTICS–AND OTHER MATERIALS–FROM PLANTS

Zachary Bogue, managing partner, Data Collective Venture Capital (USA)

For the last 80 or 90 years our innovation in materials has been driven by petroleum—by recombining petroleum compounds into fuels, plastics, drugs, and so on. I think we'll look back on the 2020s as a decade of innovation driven by biology. Genetically engineering plants to synthesize chemical compounds opens up a design space exponentially larger than petroleum, to create new materials that will let us live more sustainably and propel the economy forward. It's already starting to happen—one of the companies we invest in makes a microbe that produces a palm-oil replacement, for example. What's enabling all this is massive increases in computing power and AI that make it possible to model and design the necessary metabolic pathways.

CHINESE PHONES WILL RULE

Ronaldo Lemos, director, Institute for Technology and Society of Rio (Brazil)

By 2030 the most famous mobile-phone brands worldwide will be Chinese and they will run their own operating system, cutting the market penetration of Android in half.

"We need alternate modes of decent work—child care, health care, elder care, education."

GLOBAL SUPPLY CHAINS WILL CRUMBLE AND POOR COUNTRIES WILL SUFFER

Sharan Burrow, general secretary, International Trade Union Confederation (Australia)

3D printing, automation, and robotics will cause massive localization of manufacturing. If I can go to my local shop and I say I want my jeans with four stripes and three pockets and I want them now, the fast fashion industry is at risk. Food production will become more local too, and efforts to reduce the carbon footprint will change consumption patterns. So the supply chains on which global trade is based—dehumanizing and exploitative though they currently are—will in large part disappear from the most vulnerable countries, leaving the potential for failed states and even more desperate poverty. What we need is alternate modes of decent work, like child care, health care, elder care, education. We need to invest in human infrastructure, in support and services.

SMALL BUSINESSES WILL USE SUPERCOMPUTERS

Peter Ungaro, CEO, Cray (USA)

For example, there are hundreds of companies that make components for automotive manufacturers. Today they use small computer systems to do CAD drawings of their parts and some simulations. In future, because of all the sensors that will be out there generating data, they're going to have data sets 10, 100, 1,000 times bigger than today that they can compute on, changing how they model their parts. The technology they'll do that with will be like a mini supercomputer. Some places will have one on the premises, and others will just access it via the cloud. And it won't have to be one of these machines that today fill up two basketball courts and consume 30 megawatts. We'll have it down to a single cabinet.

WIKIMEDIA (BELL, LEMOS); FINANCIAL TIMES (BOGUE); WORLD ECONOMIC FORUM / JOLANDA FLUBACHER (BURROW); JONES / OAK RIDGE NATIONAL LABORATORY (UNGARO)

NEVER MIND THE BALLOTS

Forecasters made a lot of bad predictions during the 2016 presidential race.

On the night of November 8, 2016, Charles Franklin, like millions of other Americans, watched the presidential election results roll in with what he described as "a sinking feeling." But Franklin, a Wisconsin pollster and professor of law and public policy at Marquette University, wasn't distressed on account of his personal political preferences; he had his reputation at stake. Just a week earlier, his own poll had shown Hillary Clinton up six points in Wisconsin. Instead, here she was, losing by seven-tenths of a point.

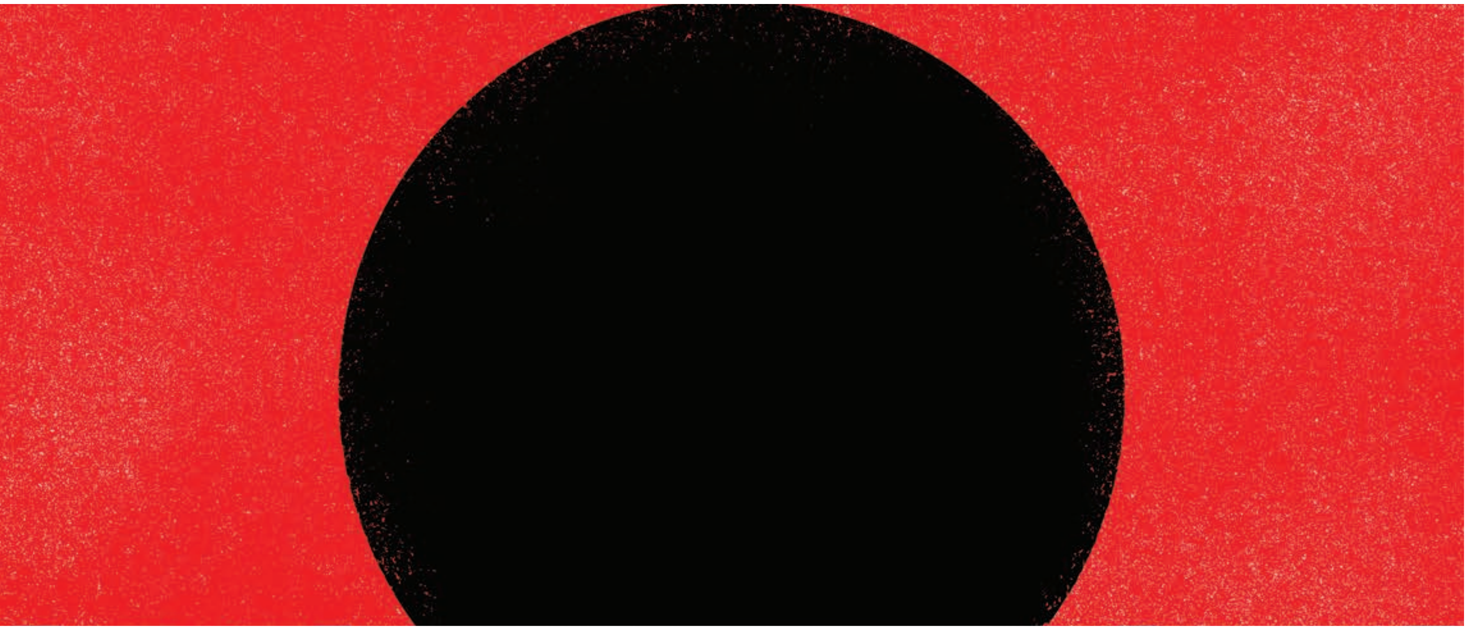

Now a crowded field is trying to get things right for 2020.

By **ROB ARTHUR**

Illustrations by Karsten Petrat

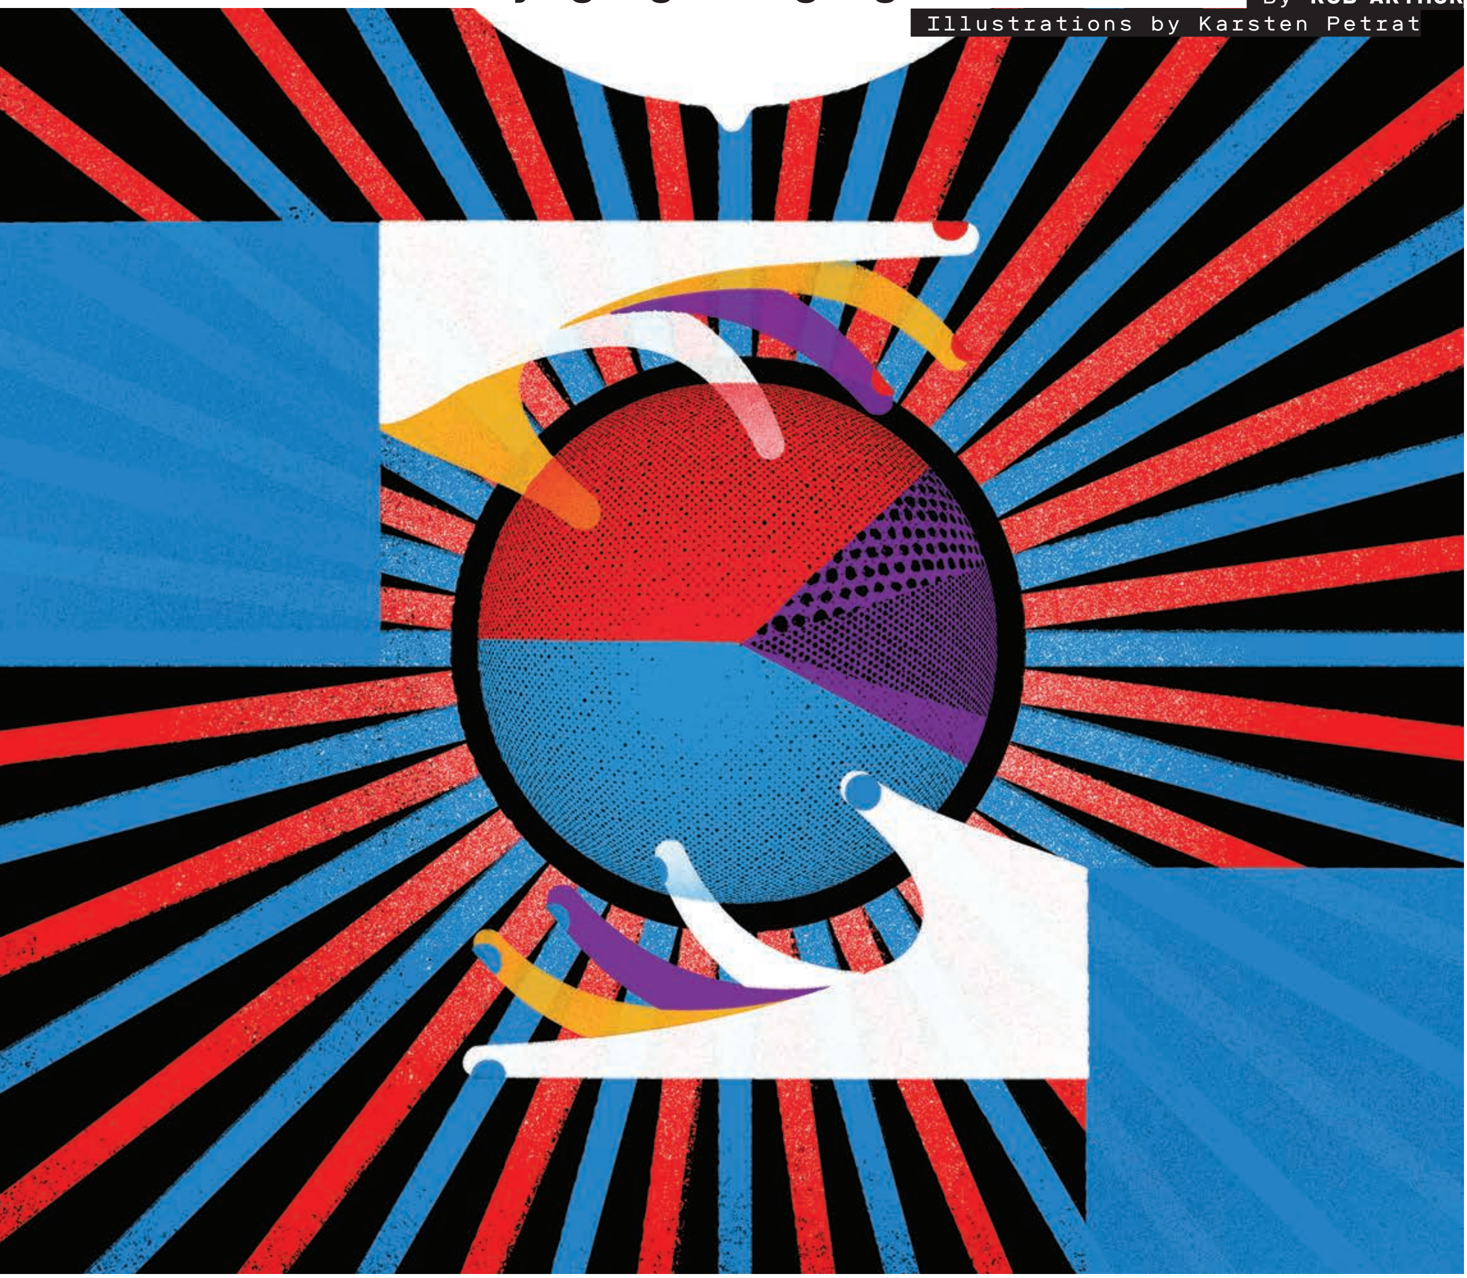

Franklin was on duty with ABC's Decision Desk, one member of an expert behind-the-scenes team responsible for calling states for Clinton or Donald Trump as the tallies came in. As he watched the returns pile up until four in the morning, it became clear that his survey was off.

"Nobody wants to be wrong," he says, looking back. "So in that sense it was very depressing."

He wasn't the only pollster to misread the election. According to RealClearPolitics, every single one of more than 30 polls in Wisconsin in the months leading to the election had Clinton winning the state by margins ranging from 2 to 16 points. And these errors had been amplified further because they were then used as fuel for computer algorithms that predicted an overall Clinton victory.

After Donald Trump had made his victory speech and the dust had cleared, everyone started to admit their errors.

"It gutted me to realize I had been wrong," wrote Natalie Jackson, a data scientist at the Huffington Post, which had given Clinton a 98% chance of winning.

The media, including many outlets whose own forecasts had given Clinton a strong likelihood of victory, started to decry the failure of prediction algorithms. Some critics were more circumspect than others, acknowledging that some forecasters had accurately described a Trump victory as merely improbable. But many cast doubt on the whole idea of predicting elections. Some even used the election as ammunition to attack the entire field of data science.

Yet nearly four years later, and with another contest looming, forecasters are beginning to issue early predictions for 2020. The backlash to 2016 hasn't dissuaded them—in fact, there's now a whole new crowd of would-be oracles, determined not to replicate the mistakes of their predecessors.

What went wrong

A cocktail of problems led to the polling misses of 2016. Some surveys failed to contact enough less-educated white voters, while some Trump supporters declined to admit which way they would be voting. Trump's unconventional strategy also turned out more citizens in heavily Republican rural counties. Pollsters incorrectly assumed that these people would stay away as they had done in previous elections, which made Trump's base appear smaller than it really was.

But while pollsters received the majority of the blame, perhaps more condemnation ought to have fallen on the forecasters, who turn pollsters' data into predictions.

"Two major forecasters had Hillary Clinton at 99% to win," says G. Elliott Morris, a data journalist at the Economist who works on election forecasting. "When she didn't, a lot of them just blamed pollsters, because it's easy for them."

There were at least two major errors committed by some of the data scientists who helped design the prediction algorithms. First, they assumed that if the odds of being off by nearly seven points in Wisconsin were low, the odds of a comparable error in other critical states like Michigan and Pennsylvania were tiny. In fact, polling problems in one state were correlated with mistakes in other, similar states. Assuming that polls were entirely independent of each other—rather than reflecting the same reactions to the same issues—produced overconfidence in Clinton's lead.

Second, prediction algorithms failed to register the record number of undecided voters as a warning sign. Because so many voters were on the fence right up to Election Day—and would end up breaking strongly for Trump—Clinton's margins were much less safe than they appeared.

"It was staring us right in the face," says Rachel Bitecofer, a professor of political science at Christopher Newport University. Had there been more polls in the closely contested states just before Election Day, she suggests, analysts might have picked up on the unusually high number of voters who decided to turn out at the last moment.

It wasn't just the forecasters' fault, though. Even when their probabilities for each candidate were accurate, the public seemed to have trouble comprehending the meaning of those numbers.

During the closing days of the election campaign, I was working at FiveThirtyEight, one of the most prominent outlets making predictions. My job didn't involve the presidential race: instead, I was covering baseball's World Series. When the Chicago Cubs were down three games to one in the seven-game series against the Cleveland Indians, I noted that their odds of winning, at around one in six, were a hair below Trump's chances of taking the White House. Six teams had done it before in the 113-year history of the World Series, and another seven had pulled it off in other playoff rounds, so it was definitely possible, but it wasn't typical. Afterwards, when both the Cubs and Trump won against the odds, I received a deluge of hate tweets blaming me for somehow jinxing into existence two very possible turns of fate.

"If you hear there's going to be a 20% chance of rain, you don't bring your umbrella. And then it rains and you get all ticked off and it's probably your fault," says Steven Shepard, an editor and election forecaster at Politico. "But that 20% occurrence isn't necessarily that unlikely."

Many people seemed to look at which candidate was projected to win (usually Clinton) without considering how certain the forecasters were. A 70% chance of a Clinton victory certainly favored the Democrat, but ought to have been viewed very differently from a 99% chance.

Still, some did say 99%, and they were undoubtedly too aggressive. Sam Wang at the Princeton Election Consortium estimated Trump's chances at less than 1%, and even pledged to eat a bug if Trump earned more than 240 electoral votes.

When the election result came through, Wang stayed true to his word. A week after polling day, he appeared on CNN with a can of "gourmet" crickets ("gourmet from the point of view of a pet," he clarified) and decried the spectacle his bet had caused. "I'm hoping that we can get back to data, and thinking thoughtfully about policy and issues," he said before dipping a cricket in honey and, with a pained expression, gulping the insect down.

Triple threat

Not all forecasts were as far off as Wang's. Some even anticipated a victory for Trump. To understand why they came in so differently, it's valuable to look at the range of approaches, which fall into three broad classes.

The earliest forecasts in each election cycle come from what are called fundamentals models. These are typically built from presidential approval ratings, economic statistics, and demographic indicators. A strong economy presages victory for the incumbent's party, as does a high approval rating for the incumbent. The demographic makeup of a state can also be used to predict the outcome—for example, white, non-college-educated voters tended to vote for Trump in 2016, so states with lots of them are more likely to go his way in 2020 as well.

Because these factors are relatively stable, reliable fundamentals predictions can be made much earlier than most other types of forecast. Models like this seem too simple to capture all the quirks and scandals

Political forecasting: The basics

Fundamentals

These models use factors like the president's approval rating and the state of the economy to predict the next winner. Many of them make predictions well in advance of the election—some are already out for 2020—but most predict only the national popular vote, and in two of the last five elections the Electoral College has swung the other way. Many were accurate to within a percentage point of the final national vote.

Examples: Alan Abramowitz's Time for Change; Rachel Bitecofer's Negative Partisanship model

Strength: Stability; long-term prediction

Weakness: Misses short-term factors; focuses on national popular vote rather than Electoral College

Quantitative

These models ingest daily polling data from both state and national surveys and rely on elaborate statistical reasoning to generate predictions. Because they are based on a rolling average of polls, they tend to fluctuate more than some other approaches, reacting sometimes to scandal-driven news cycles or shocks in the economy. The accuracy of these forecasts varied widely in 2016: FiveThirtyEight was one of the least bearish on Trump, but others gave Clinton a more than 99% chance of winning.

Examples: New York Times Upshot; FiveThirtyEight

Strength: Can react to changing circumstances

Weakness: Hard for the public to understand

Qualitative

Instead of running an algorithm, these forecasters elect to make predictions based on their own mental models of the contest. Their forecasts are updated all the way up to the day of the election and typically harvest information from all available sources, including the output of quantitative models, the state of the stock market, and the spin on the latest scandal. Some can be surprisingly accurate—on par with the best quantitative forecasts.

Examples: Larry Sabato's Crystal Ball; Cook Political Report

Strength: High accuracy

Weakness: Broad rather than specific of the modern, two-year campaign. But they performed shockingly well in 2016: six out of 10 predicted the final popular vote to within one percentage point.

The presidency isn't chosen by straight-up national popular vote, however, and that's a key limitation of fundamentals approaches: few predict the final results of the Electoral College.

Fundamentals models have another weakness. If late-breaking news arises, such as a scandal at the end of the race or a sudden shift in the economy (the 2008 financial crisis is a good example), then these stable forecasts can suddenly become woefully out of date. To compensate for this, a decade or so ago statisticians started popularizing new kinds of quantitative models that aren't quite as vulnerable to these October surprises. They process polling data as it comes out and produce a day-by-day estimate of who will win, so they can respond if public opinion shifts.

RealClearPolitics and the New York Times' Upshot both have well-regarded quantitative models, but no model has more fame—or, arguably, a better track record—than Nate Silver's FiveThirtyEight forecast, named for the total number of votes in the Electoral College. FiveThirtyEight's algorithm comes in several variations, but all take care to adjust polls according to how trustworthy the survey organization is and whether its results tend to consistently lean Democratic or Republican. The careful ingestion of polling data, and the attention Silver pays to uncertainty, have traditionally set it apart from other forecasts. "FiveThirtyEight is the gold standard," Bitecofer told me.

Of the major quantitative election predictions, FiveThirtyEight's was the most conservative, assigning Clinton a 71.4% chance to win on the eve of the election. "That sounds about right now in retrospect," says Charles Franklin: Trump's victory was an unlikely, but not impossible, outcome.

Finally, there are predictors out there who eschew mathematical approaches altogether, relying instead upon a combination of intuition, polling, and the output from all the other kinds of models put together. These qualitative predictions run on one of the most sophisticated and yet error-prone computational engines we know of: the human brain.

Rather than precise numeric estimates, qualitative forecasters typically group races into one of four categories on a scale ranging from safe to toss-up.

"Toss-up" means there is no favorite: "Kind of a coin flip," says Kyle Kondik, a qualitative forecaster with the University of Virginia's Crystal Ball political analysis newsletter. "Lean," he says, is a small edge for one side or the other. "Likely" is a larger edge for one side or the other. And "safe," he says, means we'd be shocked if that party lost. Some

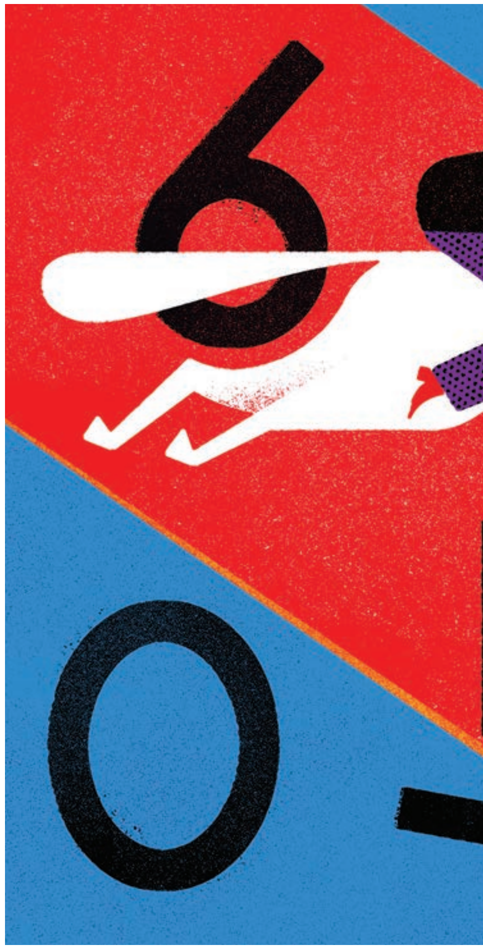

On target: Who got it right and wrong in 2016

Forecaster	Type of forecast	Final prediction
Nate Silver/ FiveThirtyEight	Quantitative	71.4% chance of Clinton victory
Sam Wang/ Princeton Election Consortium	Quantitative	93-99% chance of Clinton victory
Cook Political Report	Qualitative	278 electoral votes for Clinton (46 of them toss-up votes)
Larry Sabato's Crystal Ball	Qualitative	322 electoral votes for Clinton
Alan Abramowitz's Time for Change model	Fundamentals	51.4% of the national vote for Trump
James Campbell's Convention Bump model	Fundamentals	51.2% of the national vote for Clinton

“HORSE RACE POLLING IS BELIEVED TO INCREASE CYNICISM ... IT CAUSES PEOPLE TO VIEW POLITICS AS A GAME, WHERE THEY GO OUT AND ROOT FOR THEIR TEAM.”

qualitative predictors argue that these verbal groupings help readers understand the relative probabilities better than the more exact numbers offered elsewhere.

While these predictions may seem less scientific than ones based on crunching numbers, some boast an impressive level of accuracy. In the 2018 midterms, according to a third-party assessment of several professional forecasts, it was the aptly named Crystal Ball that did best, not FiveThirtyEight’s statistical algorithm. Performance tends to fluctuate from cycle to cycle, however: the best practice, according to pollsters and academics, is to consume a wide variety of forecasts—qualitative, quantitative, and fundamentals.

What next?

Nearly all the forecasters I spoke to had received vitriolic hate mail after the 2016 results. Yet dozens of new modelers have thrown their hats into the ring for 2020.

They will be rolling out their predictions for the first time this year, and they are intent on avoiding mistakes from past election cycles. Morris, the Economist’s forecaster, is one of those entering the field. He has called previous, error-prone predictions “lying to people” and “editorial malpractice.” “We should learn from that,” he says.

The Economist will be building its algorithm using polls published by outside organizations, but it will also be conducting its own surveys to shore up the results in ambiguous states and races, which Morris hopes can lead to greater accuracy.

The Washington Post, too, is making its first gamble on predictions—but taking a different route. It is staying out of the forecasting game until returns start coming in. Only once the first precincts start to announce vote totals on Election Day will the Post deploy its analytical model to judge the likelihood that specific candidates take the state or district for which they are competing. By waiting until the first ballots are counted, the Post’s data scientists plan to drastically reduce the error in predicting the rest of the votes, albeit at the cost of being unable to release an early projection.

Experienced forecasters and pollsters aren’t sitting on their hands either. Builders of fundamentals models are beginning to take up the challenge of predicting the Electoral College instead of just the popular vote. Bitecofer designed a model based primarily on demographics that is already predicting a narrow electoral-vote victory for the Democratic challenger, whoever that may be.

The designers of those problematic quantitative algorithms appear to have learned their lesson about correlated errors between states. The Huffington Post issued a *mea culpa* for its 98% prediction of a Clinton victory. Wang, the bug-eating Princeton professor, has pledged to update his algorithm so that it will be much less confident in 2020, admitting on his blog that his earlier model was “a mistake.”

Qualitative forecasters, meanwhile, took a variety of lessons from 2016. “There are a lot of different things that in hindsight I wish that maybe we had focused on a little bit more, but I would say the fundamentals-based models were the best in that election," says the University of Virginia's Kondik. "I wish we all paid them greater heed."

Kondik and others stress the need to be cautious about any prediction given the historic unpopularity of the sitting president, which ought to decrease his chances, and the strong economy, which ought to increase them. Those dueling factors mean the race is uncertain so far from Election Day.

Elsewhere, media organizations have also started providing their estimates in ways that are designed to give the reader a better, more intuitive grasp of what probabilities mean. Rather than writing that Democrats had an 87.9% chance of taking the House during the 2018 midterm elections, for example, FiveThirtyEight emphasized that they could have expected to win seven times out of eight.

"Psychologists have found that people are better at understanding these types of [numbers]," wrote Yphtach Lelkes, a professor of communications at the University of Pennsylvania.

Finally, pollsters are upping their game as well. The American Association for Public Opinion Research (AAPOR) issued a retrospective of 2016 with lessons for future elections. Tips include using statistical tricks to ensure that population samples are more representative of the state being surveyed and conducting more polls in the final days of the campaign so as to capture the leanings of late-deciding voters, who proved so critical to Trump's victory.

Franklin, the Wisconsin pollster, was one of the authors of AAPOR's post-mortem. The systematic failure of dozens of surveys across several states suggest that his poll's mistake was due to a real shift in the closing days of the race, rather than an earlier, more fundamental error. Still, he wonders what might have been: "What if we had polled through the weekend before the election? Would we have captured the swing toward Trump in those data?"

Quantum polling

But while mistakes from four years ago can be corrected, new difficulties may also crop up for the 2020 cycle. Some may even be driven by forecasting itself. Some experts argue that election predictions may be influencing the very results they are trying to predict.

According to a recent study, an overwhelmingly liberal audience tuned in to those overly confident quantitative forecasts in 2016. Previously published studies suggest that when people believe the outcome of an election is certain, they are less likely to vote, especially if the certainty is stacked in favor of their chosen candidate. So in a twist on what is known as the observer effect—in which the mere act of watching something changes the outcome—feeding a heavily Democratic audience with a steady diet of overconfident polling like Wang's could have reduced turnout significantly. Given that the race was essentially decided by only 107,000 votes in three states, any reduction could have been important.

STUDIES SUGGEST THAT WHEN PEOPLE BELIEVE THE OUTCOME OF AN ELECTION IS CERTAIN, THEY ARE LESS LIKELY TO VOTE.

"Clinton lost by so few votes that it is certainly possible that probabilistic forecasts caused enough Democrats to stay home that it affected the outcome," wrote Lelkes. Clinton herself suggested as much. "I don't know how we'll ever calculate how many people thought it was in the bag, because the percentages kept being thrown at people—'Oh, she has an 88 percent chance to win!'" she said in an interview in New York magazine.

Even if election forecasting didn't change the outcome in 2016, it could have more of an impact on future campaigns.

"Horse race polling is believed to increase political cynicism, affect turnout, increase polarization, and likely supplants information about substantive issues," wrote Lelkes. "It causes people to view politics as a game, where they go out and root for their team, rather than support candidates based on their political positions." And if these effects are real, they are likely to get more powerful as more forecasts happen.

Some forecasters, like Silver, have dismissed this concern. They argue that it isn't their job to tell people whether or not to vote—or to tell the media what to cover. Others, however, are taking the advice of Lelkes and his colleagues more seriously.

"We're experimenting with ways to convey uncertainty that won't turn people off [from voting]," says the Economist's Morris. "But I think that is still a problem that forecasters are going to have … I don't know how we get around some of the societal implications of our work."

Rob Arthur is an independent journalist and data science consultant based in Chicago.


Wharton
UNIVERSITY of PENNSYLVANIA
Aresty Institute of Executive Education
EXECUTIVE EDUCATION
excel
verb \ik·ˈsel\
"The moment I realized I had become part of something that would take me further than I've ever been."

5 of the best books about prediction

Predictions of any importance are never only about saying what will happen. Right or wrong, they also shape the course of events. The types of predictions that are possible keep changing: the present is modulated by beliefs about how much is knowable about the future. These books discuss consequential shifts in how predictions are conceived.

By **KONSTANTIN KAKAES**

Moneyball

MICHAEL LEWIS

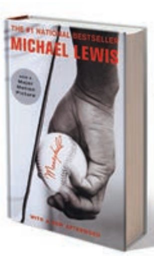

This 2004 book about how the Oakland A's set out to use better metrics to predict baseball players' skill has become a landmark not just in sports writing but in how "data-driven" decisions became popular in all walks of life in the 21st century.

"Pitchers were like writers in another way, too: their output was harder than it should have been to predict. A twenty-two-year-old phenom with superior command wakes up one morning in such a precarious mental state that he's hurling pitches over the catcher's head. Great prospects flame out, sleepers become stars. A thirty-year-old mediocrity develops a new pitch and becomes, overnight, an ace. There are pitchers whose major league statistics are much better than their minor league ones. How did that happen? It was an odd business …"

Surfing Uncertainty: Prediction, Action and the Embodied Mind

ANDY CLARK

Philosopher Andy Clark argues that prediction is central to understanding how human beings perceive the world. We do, he says, a lot more predicting than we realize.

> "Prediction is, of course, a slippery beast. It appears, even within these pages, in many subtly (and not-so-subtly) different forms. Prediction, in its most familiar incarnation, is something that a *person* engages in, with a view to anticipating the shape of future events. Such predictions are informed, conscious guesses, usually made well in advance, generated by forward-looking agents in the service of their plans and projects. But that kind of prediction, that kind of conscious guessing, is not the kind that lies at the heart of the story … Brains like ours, this picture suggests, are predictive engines, constantly trying to guess at the structure and shape of the incoming sensory array. Such brains are incessantly pro-active, restlessly seeking to generate the sensory data for themselves using the incoming signal (in a surprising inversion of much traditional wisdom) mostly as a means of checking and correcting their best top-down guessing."

Cassandra

CHRISTA WOLF

In this novel, Wolf tells the story of Cassandra, a daughter of King Priam of Troy, who was able to predict the future, but whose predictions were never believed.

"I alone saw. Or did I really 'see'? What was it, then? I felt. Experienced—yes, that's the word. For it was, it is, an experience when I 'see,' when I 'saw.' Saw that the outcome of this hour was our destruction. Time stood still, I would not wish that on anyone. And the cold of the grave. The ultimate estrangement from myself and from everyone. That is how it seemed. Until finally the dreadful torment took the form of a voice; forced its way out of me, through me, dismembering me as it went; and set itself free. A whistling little voice, whistling at the end of its rope, that makes my blood run cold and my hair stand on end."

Gravity's Rainbow

THOMAS PYNCHON

One of the 20th century's most important novels, Gravity's Rainbow, *among other things, tells the story of Tyrone Slothrop, who can predict where German V-2 rockets will impact in London. Researchers in the novel ponder whether Slothrop is causing the rockets to fall where they do, or merely predicting.*

"The rockets are distributing about London just as Poisson's equation in the textbooks predicts. As the data keep coming in, Roger looks more and more like a prophet. Psi Section people stare after him in the hallways. It's not precognition, he wants to make an announcement in the cafeteria or something . . . have I ever pretended to be anything I'm not? all I'm doing is plugging numbers into a well-known equation, you can look it up in the book and do it yourself. . . . His little bureau is dominated now by a glimmering map, a window into another landscape than winter Sussex, written names and spidering streets, an ink ghost of London, ruled off into 576 squares, a quarter square kilometer each. Rocket strikes are represented by red circles. The Poisson equation will tell, for a number of total hits arbitrarily chosen, how many squares will get none, how many one, two, three, and so on. An Erlenmeyer flask bubbles on the ring. Blue light goes rattling, reknotting through the seedflow inside the glass. Ancient tatty textbooks and mathematical papers lie scattered about on desk and floor."

The Taming of Chance

IAN HACKING

This brief, dense, and beautiful work of philosophical history explains how, in the late 19th century, statistics acquired explanatory power for the first time as scientists gradually abandoned the idea that the present completely determines the future.

"The bad player is the one who tries to calculate and play with the odds, as if his game, his life, were one of a large number of games. To do so is at best to succumb to another necessity, the necessity of the law of large numbers. The good player does not fool himself, and accepts that there is exactly one chance, which produces by chance the necessity and even the purpose that he experiences. Not even a long run of universes would annul the chance that brought into being our world, and only the false consciousness of a bad gambler could make it seem otherwise."

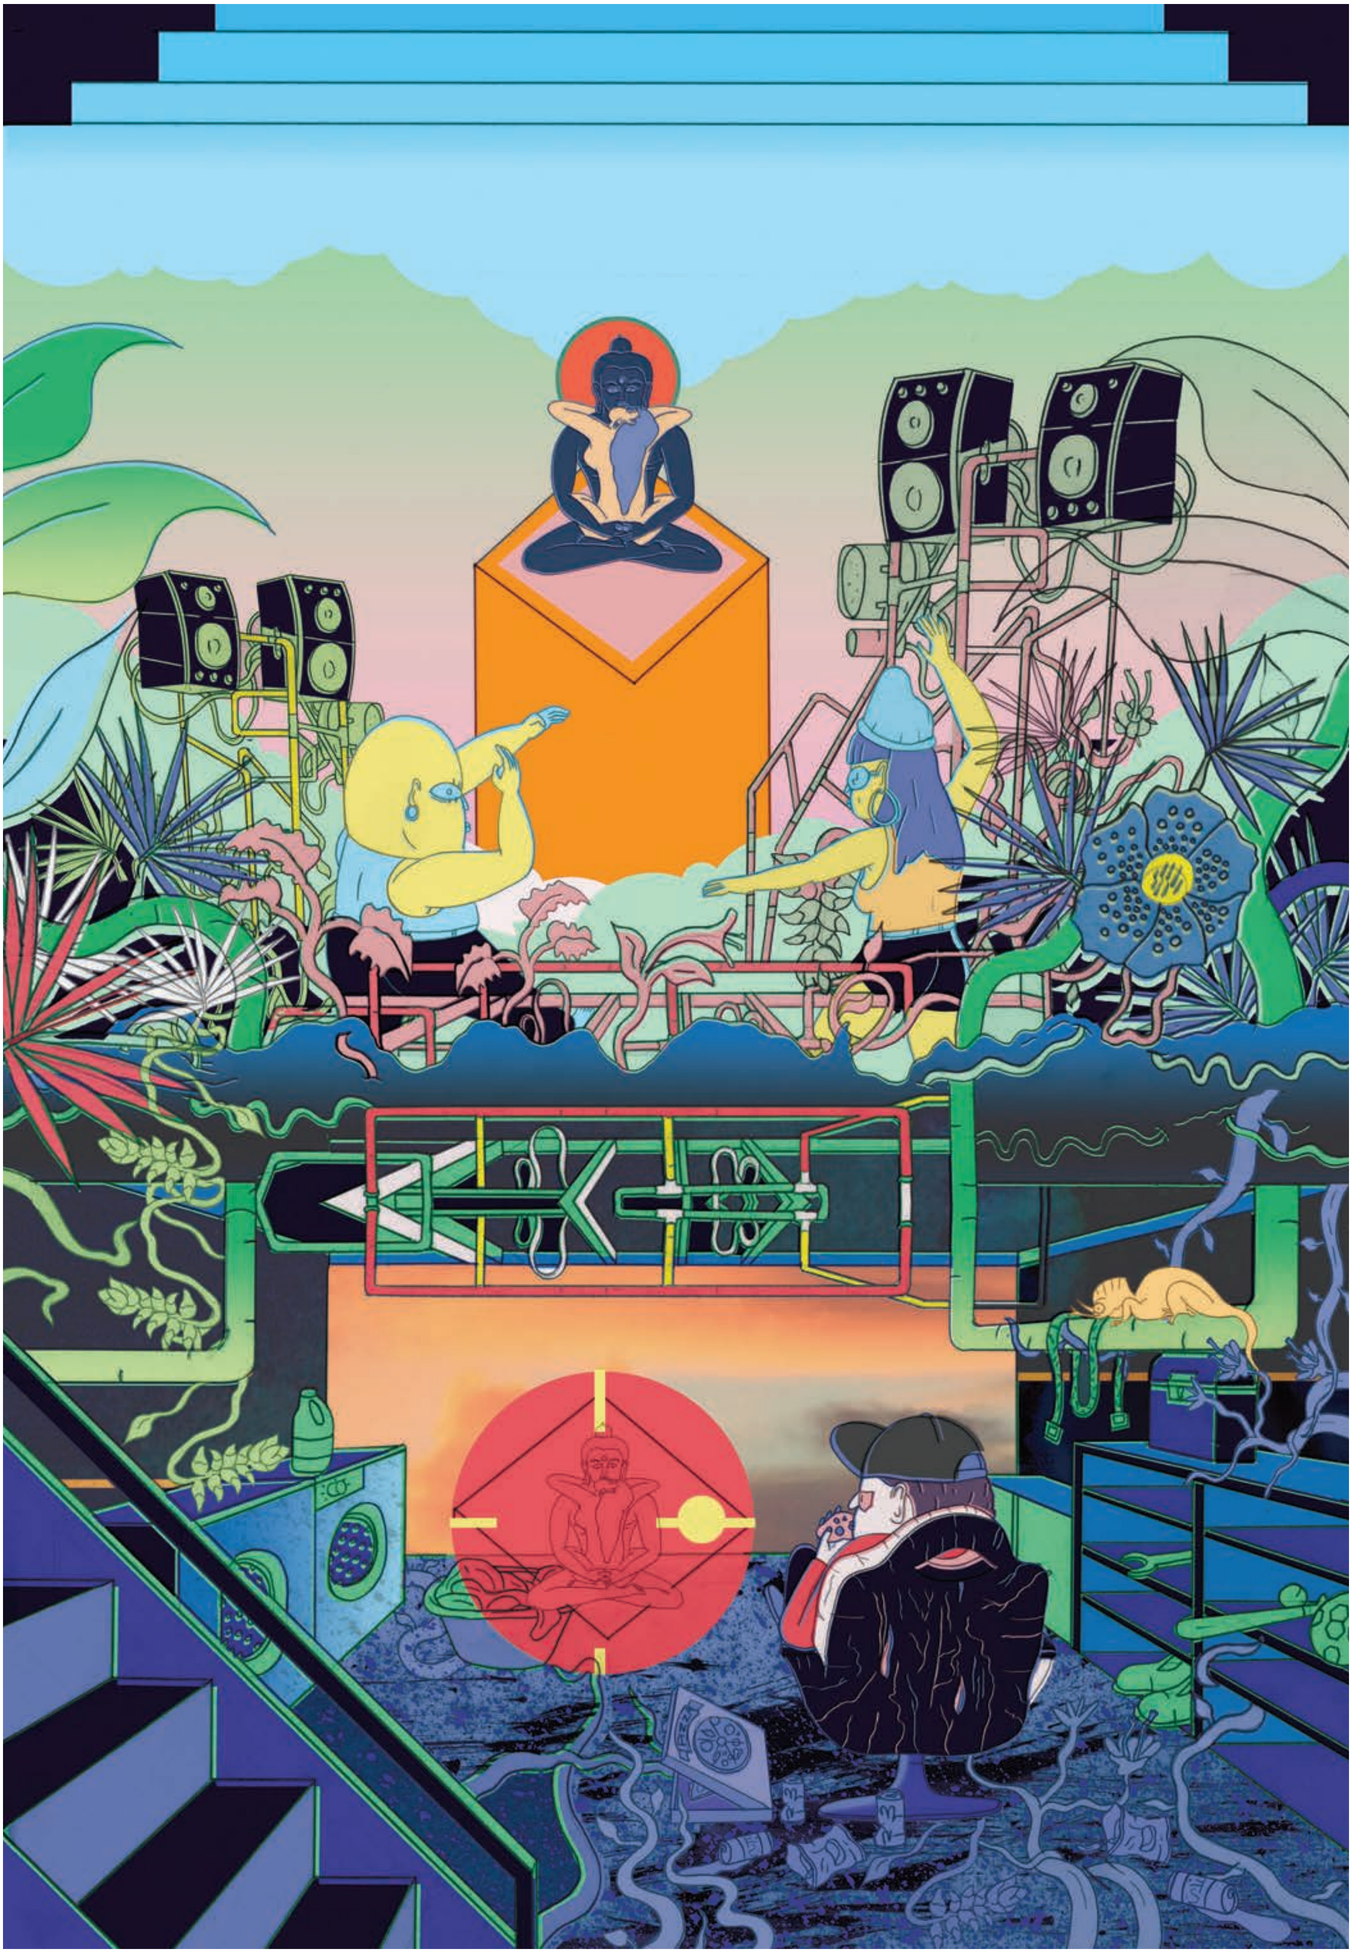

Fiction

Zooming

I'm sitting in my parents' basement, in a cracked pleather gaming chair, smelling my own funk, or maybe the damp of black mold, and 400 miles below me the whole world is laid out like some vast Tibetan tapestry, full of little demons and beasts and believers.

I tap, zoom, look, unzoom, slide, tap, zoom, look. Sometimes at familiar spots, but mostly just at random, searching for something happening somewhere that's interesting enough to stream or gif or sell or just linger over. I watch Berliners mob a music festival. I watch mining equipment drag rocks out of an Australian quarry. I watch Pakistani dogs fighting over a chicken and hurricane clouds slamming into Cuba and an exhibitionist couple fucking on a bright red blanket on a Californian rooftop. I lose myself for a few minutes in the ripples of swaying Amazon jungle leaves, wondering how the wind feels to all those trees. And then I get bored, and I'm just zooming through my rounds again, not thinking much, and I see it.

Some kid is dragging a tasteful brown coffin out of the back of a pickup truck parked at the edge of a pile of trash in the junkyard just outside of town, my town. Silent *thunk* when the box hits the trashdirt, and the kid loses his grip, rolls it, and out comes a body. Denny's body.

Never seen him from this angle before, fat face sprawled to the open sky, but somehow I know it's him: the lima bean bald spot who wore a hideous Hawaiian shirt on their first date, just like the body is wearing now. Denny is the guy fucking my ex Michelle. Was the guy, because I'm pretty sure I'm looking at a live satellite feed of his corpse.

I zoom as hard as I can, but the algo caps the resolution when it thinks there are people in the frame. Panoram doesn't want us swiping credit card numbers or peeking at text messages, even though they probably sell that data to marketing firms or use it to blackmail Saudi princes. I can see the coloration on individual feathers on a bird soaring over some pristine wilderness, but trying to identify a dead body is like spotting an acquaintance across

BY ANDREW DANA HUDSON
ILLUSTRATIONS BY DOGBOY

Too much of life happens inside, underground, in cars or trains, under trees, on cloudy days.

the street through a smudgy bus window. Doesn't matter how sure I am—no one else will believe me.

The kid plants his hands on his hips for a minute, then bends to shove Denny back in the coffin. He gets the lid on, latches it, I guess, and gives the coffin a couple rolls toward the junk pile.

I don't do snuff zooms, even though they're good money on the dark web. I don't chase car crashes or predator drones or active shooters. I should bug out, look at something else, watch a nudist beach or contemplate some cracking, melting ice floe. Everyone knows Panoram can't afford storage for all the imagery it takes, if storing that much data is even possible. If a user doesn't record it, it's gone forever—the tech-god is omniscient but forgetful. I could pretend I never saw Denny's blurry pixel eyes staring up at me.

But death is weird when it's someone you know, even if they didn't know you. I never met Denny in person. I only know his name from my buddy Trent who still goes to Michelle's restaurant sometimes. Still, I've watched Denny pick Michelle up from barre class, drop her off at work the next day. Little flick of the wrist as he called her back for one last kiss. Maybe I was jealous, but I didn't hate him. We shared a world, and now someone's thrown him dead in the garbage.

So I hit Record. Seems like the least I can do.

The kid wipes his brow, like "Another day, another dollar," and I'm sweating just looking at him, itching at my pits, peering desperately into my monitor for some detail on the kid beyond the slightness of his frame and his logo-less baseball cap and grubby black T-shirt. But there's nothing. Kid gets back in the pickup. It drives off.

I zoom out to follow. Long shot, but who knows where amateur body-dumpers get their vehicles. Couple miles from the junkyard, the truck turns in to a covered garage where empty fleet cars go to charge. I circle around the shiny black square of solar roof for a few minutes, just in case the kid hoofs it. Windowless sedans zip out of the hub like blind ants, leaving their anthill on pheromonic marching orders. He's probably already in one, napping off the sun. I've lost him.

But I do have a time stamp. Silver pickup entered the hub at 11:28:15 MT. Just like in crime shows, the cops can warrant the garage logs, track the truck back to wherever it picked up the kid—and Denny's coffin.

I should ping the cops. But I don't, because there's something else I've seen in crime shows. One in five homicides are committed by an intimate partner, which means there's a non-zero possibility that Michelle was the one who had Denny offed. What if he beat her? Or stole her money? Or tried to sexually traffic her? I'm a snitch, but I'm not going to snitch on her.

My best bet is to find Michelle, keep recording the evidence, track her until I get the whole, fatal story. I pull an Adderall shot from my minifridge, slosh it down, toss the little can, purple liquid splatter joining the salsa stains on the wood-grain carpet. I order pizza to the basement door, text Mom and Dad that I'm staying in. It'll be at least a day before they throttle my bandwidth to force me upstairs. I go to the bathroom and scrub caffeine on my face. Then I go looking for Michelle.

The thing about zooming is, it's actually fucking hard to stalk people. Too much of life happens inside, underground, in cars or trains, under trees, on cloudy days. And they know we're watching, so floppy hats are back in a big way, gated communities put up shade sails, couples kiss under umbrellas on rainless afternoons.

Then there are the anti-stalking algos that kick you off if you zoom in on the same address too long or too often. Panoram is for wildlife photography and storm chasing and seeing humanity in its broadest strokes: the daily heaving of commuters, migrants, pilgrims, supply chains, shipping lanes, air travel, construction sites, battle lines, strip-mining, clear-cutting, controlled burns, cook fires, city lights, parades, sports games, mass weddings, protests, riots.

Finding Michelle is like finding a needle in a haystack when the haystack is on fire. Impossible—except I've had a lot of practice.

I catch her coming out of the Thai place when her shift ends after the lunchtime rush. I know it's her from the way she twists her hair up into a bun and the stretch she does, there on the sidewalk, to celebrate being off the clock. She's unbuttoned her white hostess shirt, down to a sweaty halter top, and the slight angle of the satellite lets me gaze right into her pixelated cleavage. She arches her back like she wants me to see.

Everyone checks up on their exes, right? I don't want her back, but I zoom her when I want a reminder that she's hot, cool, and successful, and for a while she chose me. Or else I want evidence that she's miserable and pathetic without me. Or maybe she's ugly, tacky, slutty, immoral, and I'm better off without her, better than her, now that

I've come to my senses and moved on. Or none of that. It's just an itch to scratch.

Today she's got a bounce in her step, like she got a really good night's sleep or maybe got away with murder. She's not checking her phone or edging away from passersby or any of the nervous movements I'd expect from someone whose boyfriend has gone missing, who's involved in a criminal conspiracy, who's about to go on the lam.

Michelle walks to the library, comes out 10 minutes later. She goes to a coffee shop, spends an hour inside. To keep the algo from getting suspicious, I pan over the café slowly, jump to a random spot, then come back and sweep the surrounding blocks in case I missed her. Rinse, repeat. My pizza arrives. It's pure luck that I catch her leaving.

More errands. I haven't zoomed on one person this long since I watched a Mongolian nomad track a runaway horse two days across the steppe. I've followed Michelle before, but always with a bored, idle, compulsive curiosity—never with actual focus.

She goes to barre class. I figure this is it. When she's done, either she'll wait for Denny to pick her up until she realizes he's not coming, or she'll just go, because she already knows where Denny is.

Fifty minutes later the studio empties. A dozen pairs of yoga pants come out, all buzzing with post-workout endorphins. They scatter, but not Michelle. She waves them off, plops down on the curb, waits.

I get this rush of relief, and I'm about to call the cops, tell them about Denny—anonymized so there are no questions about why the victim's girlfriend's ex-boyfriend knows where the body is—when a car pulls up.

From my vantage, it's a windowless black lozenge. A side panel opens, and out leans the same black T-shirt and cap, same slight arms that rolled Denny onto the trash heap this morning.

I want to scream down from the heavens, blare on some global satellite PA system, warn her: *Do not get in that fucking car.*

She gets in the car. It drives off.

It's rush hour now, and tracking the car is like playing Grand Theft Auto and Frogger and a street hustler's shell game. I ache for the days of early Panoram, when they still let in third-party algos that could track vehicles and individuals for you. Dozens of identical sedans merge and exit in a tight, automated gridlock, and I go cross-eyed trying to stare at the one Michelle is in.

Either my ex is heading off into the sunset with the hit man she hired to get rid of Denny, or she's riding around with a killer and has no clue how much danger she's in.

I call her phone. No answer. I text her: *Jump out of that car!* That gets her attention. She calls me.

"Shawn, you can't keep doing this," she says. "I deserve privacy—you agreed! If you zoom me again, I'll ... I'll report you to Panoram. I'll get a restraining order."

I tell her it's not like that. I tell her she's in danger. I tell her I saw the guy in the car dump the body.

She says, "What body?"

So I tell her to open Panoram on her phone and zoom on the trash pile in the junkyard just outside of town, our town. I ping her the coordinates and tell her to look for a coffin.

Pause with some heavy sighs as I guess she does what I ask. Then: "I don't see anything but garbage and big crane things."

Dogboy, also known as Philip Huntington, is a London-based illustrator working toward the illustration of an alternate reality. His ongoing project Dystopolis is the fictional documentation of a dystopian society.

My fingers twitch and pinch, and with a bolt of shame, I realize I want to zoom on the box.

I zoom back to the junkyard on my own screen. A pair of earthmovers are rearranging the trash pile right where Denny's coffin had been. Fuck.

I tell her she has to believe me.

She says, "Shawn, how long have you been staring at that screen? Maybe you should get out."

Fine, I say. Fine. I'll show you. I send her my location. Then I get out of my chair.

In the garage is the bike I never ride. My dad keeps the tires pumped up because he read a book about how the best way to parent my generation is to remove the obstacles that prevent us from exiting self-destructive behavior. I clip in my phone, roll out of the garage, immediately start sweating in the sunset heat.

Riding the bike again is just like riding a bike, but harder. My legs ache, my lungs burn. I look up over my shoulder, and I try not to imagine how my soaked back, hunched over the handlebars, must look to Michelle through the satellites above.

I take the bike paths that tendril out of town—faster than rush hour traffic, even at my huffing pace.

All the while, I'm on the phone with her, trying to explain, though I'm out of breath. Eventually she says, "Okay, let me come meet you. We can figure this out." Then neither of us talks much. For some reason, I feel better, even though I know that if she is a killer, she's probably only coming to kill me too. I keep my eyes on the road, and on the blip of my body that Panoram keeps centered on the map it lays over the feed on my phone.

There's no guard at the junkyard, just a gate where you insert your credit card. All the junk is chipped, and you pay by the pound. I dismount and walk into the stacks of objects too toxic to compost, too complex to recycle, too useless to repair. After a day of looking down, their three dimensions weird me out; their perfect resolution sets my teeth on edge.

The automated earthmovers have wandered off, but I see the work they've done. They've lifted Denny's heap and set it precariously on top of an adjacent pile, a steep little hill of things no one wants. I see the brown corner of the coffin near the top, covered by a tangle of broken clothes hangers and old halogen lamps.

My fingers twitch and pinch, and with a bolt of shame, I realize I want to zoom on that box. But I can't. Instead I walk up to the hill, get purchase on a torn-open-mattress spring, and begin to climb.

The sun trickles away, and inch by rattling inch I edge up the mound of trash, toward the sky. I'm almost to the box when I hear Michelle's voice.

"Shawn! Please! You have to come down from up there!"

I crane my neck, and she's there, just how I remembered: overbleached barrel-collar shirt and sensible flats. She clutches her phone, and I can see Panoram's darkening view of the junkyard between her white knuckles. Her face is a picture of concern.

Next to her stands a skinny guy, the kid, maybe, though in the flesh he looks older. Is he angry? Stoic? Sympathetic? Territorial? I can't read him. T-shirt more green than dark, and he's ditched the baseball cap. But he's still the kid I saw, I know it, he's got to be. Except—there's this bald spot that licks over his scalp, shaped like a lima bean.

I ask who's that.

"Shawn, this is my partner Denny," Michelle says. "He came with me because he's worried. We all are. We don't want you to hurt yourself."

I tell her that's bullshit. I tell her Denny's dead.

"Shawn, come down here. Talk to us. Look me in the eye for once."

I keep climbing. I get to the coffin. From here it's not so sleek. No $10,000 polished mahogany, just stained plywood, glued together. More of a shipping box than a proper casket.

I try to tug it out of the pile. The junk shifts, but doesn't budge.

I hear whispering from below, then feel a creak. New Denny is on the pile with me, climbing.

I'm a sitting duck. Whoever this guy is, he knows I know too much. I could kick at his face, but my legs are sore from biking, cramped from sitting all day. Instead I edge away around the peak of the pile. He can't see me, but I can't see him. I pull out my phone and watch through Panoram as his bald spot picks its way up the hill.

He's going to beat me and strangle me, and then he'll probably have to kill Michelle too, bury both of us in this trash heap with his first victim. I can see it all in my head, from a god's-eye view. The way he'll put his hands on his hips after he shoves us into the garbage, wipe his brow, walk back and get a car, slip into the pool of anonymous everyones, safe from the eyes above. Our one chance at justice would be another zoomer, recording in Panoram, but what are the chances lightning will strike twice? There's no one, because no one cares about this place or this body or Michelle or me except me.

He's almost around the corner. My eyes don't leave the screen, but my free hand closes on something long and thin—one of the lamps—and I swing out to the right. The lamp rattles my arm as it hits, and I look over to see New Denny grimace, go blank, and topple. There's a moment of thick, curdled time as he falls, but then he's rolling down the pile with clank and crunch. He comes to rest rag-doll limp at the bottom of the junk heap, skinny face sprawled to the open sky.

Michelle runs forward. She screams. She's got her hands on his head and she's wobbling it, trying to make it sit right on his neck. But it won't.

I stagger down the pile. The guy lies still, except for Michelle's jostling. She's pounding on his empty chest, saying, "Shit, we shouldn't have come. Shit."

I don't feel anything, just Adderall crash mixing with adrenaline rush and cyclist high. I should go to her, comfort her, put my arms around her, but my eyes keep tugging away to the glow of the phone she's dropped. On the sepia-shifted screen I see the whole scene playing out in miniature. The blur of a woman, crouched by the blur of a body. And me, standing over them, the blur of a killer.

I pick up the phone. Panoram's red recording dot blinks at me. I know what I'd think if I were zooming this right now. I wouldn't understand at all.

I put her phone in my back pocket, squeezed next to my own, then scramble back up the pile. I get on top of the coffin, clear off the junk, and then shove. In jerks and tips, I haul the box to the ground.

Michelle is staring at me, and I don't understand her expression. She's picked up a broken chair leg from the pile, holds it at her side like a club.

"Give me my phone," she says. "I'm going to call the police. We'll tell them you had an episode, you got confused. I'll make them understand."

She doesn't know I saved her. I tell her she has to see this. I bend to work the latches.

Doubt comes to me then. For a blink, I'm expecting to find a mannequin, some haunted house prop, thrown away by a carnival, blurred by Panoram, interpreted by my brain as a vast conspiracy that I was uniquely qualified to untangle. What if there's nothing in there except my own ego, pattern recognition, and the follies of know-nothing omniscience?

But in the box there is a body.

Hawaiian shirt and a placid, pale, lumpy face. It sits at the edge of the heap, parallel to New Denny, both missing that vital force that makes meat mean something.

"Who the fuck is that?" Michelle says. She pauses, then adds, "Shawn, what the fuck did you do?"

That guy did it, I tell her. I saw it. Just zooming around, and I saw it. She should have just gotten out of the car, and I could have shown her alone, but she brought him, and he was going to kill us both.

She's shaking her head, red wet eyes full of hate and pity.

I tell her I'll prove it. I look down, dig for my phone, and she hits me. I'm on the ground, wind knocked out of me, pain screaming in my skull. I feel the two phones tug out of my back pocket. Then I get a little air, and close my eyes.

When I come to, Michelle is gone. The sun is gone too, the pink drained from the sky. The bodies are still there, but there's no hiding them now.

I stagger to the junkyard exit. Michelle has taken my bike, or someone has. I stare down the road, thinking of the silver pickup, trying to remember how far it was to that charging structure, trying to figure out if I could hoof it.

Red and blue lights start to flash in the distance. Whatever I did or didn't see, it hardly matters now. Maybe Michelle is the killer, but she has my phone, probably remembers my passcode. She can delete my Panoram recording, pin both bodies on me. Or maybe she's not, and I killed that man for nothing. Either way, when the cops get here, I'll be jailed or committed, tucked in a tiny cell with no windows, nothing to see.

I run.

I flee the junkyard and the country road, staggering through brownfields and scrubby desert until the light pollution dims to a yellow haze. Above me, the stars grow brighter, and closer. Closer still are the winking eyes of Panoram, in an endless parade of overlapping rings—satellites dancing into new constellations, filling the firmament with heroes and gods and heretics.

The police will be watching me through them. They'll have a picture-perfect view—crisp night vision, infrared. I can feel their gaze pressing on me, seeing everything about me but understanding nothing. I look for cover, but there is none. I'm exposed to the seeing sky.

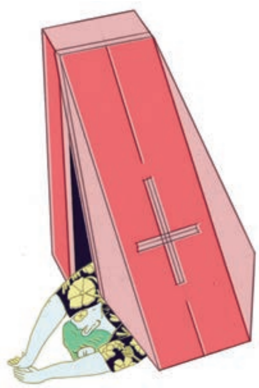

Andrew Dana Hudson is a speculative fiction writer and graduate student at Arizona State University, where he researches climate politics and AI.

A brief history of the future

This publication has been predicting what comes next for 120 years. Here's what some of those predictions have looked like through the decades.

January 1932

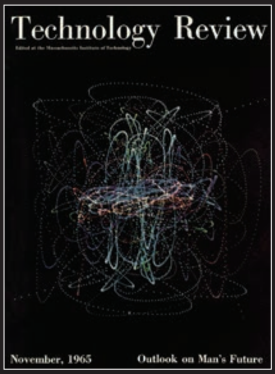

November 1965

January 1991

From "The Future of Engineering": Interpretation of the events in the past is the only valid method of predicting the future. The tremendous advances of our own civilization during the past 100 years are due more than anything else to the harnessing of the power of steam and the utilization of the energy in coal and water power. The socially important aspect of the machine age is not the machines themselves but is to be found in the fact that increased productive power has released human thought and energy for development in directions other than mere existence. If "the pathway to the future is in the hands of the engineer," then he has a very grave responsibility! It is not enough that he should produce new implements for man's desires; he must take a leading part in seeing to it that the new world which he is creating is a good sort of a place in which to live.

From "The Economics of Year 2001": The upper classes in the United States are rich, but the average income recipient has unsatisfied wants and there are substantial pockets of poverty. If we discuss 2001, however, we can pass the problem of abundance. At 3 per cent rate of growth, income per capita will be more than doubled by that year. The less developed countries may not have come very far by 2001, but they will have controlled their population rates of growth, I venture to predict, and their rates of growth in goods and services will be substantial. I assert that economics in 2001 will have the answers to a lot of questions we cannot answer today. Economics is like meteorology—a field where small differences in the relations among variables are critical to the character of the outcome. Unhappily, in its strong subjective element it has one disadvantage over meteorology. Unlike raindrops, people adjust their behavior.

From "Building the Information Marketplace": The vision I have is of an information infrastructure that would make it easy for the computers in every home, office, school, and factory to interconnect. Text, movies, software, and more would move easily over this substrate. By speeding up many of today's tasks and making possible an almost unlimited number of new activities, this infrastructure should improve our economy and our way of life. The National Information Infrastructure would make possible a United States where business mail would routinely reach its destination in five seconds instead of five days; where goods would be ordered and paid for electronically; where a retired engineer in Florida could teach high school algebra to students in New York City; where a parent could deliver office work to a distant employer while taking care of young children at home—and on the list goes.

MIT Technology Review (ISSN 1099-274X), March/April 2020 issue, Reg. US Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2020. The editors seek diverse views, and authors' opinions do not represent the official policies of their institutions or those of MIT. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: Send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the internet at www.technologyreview.com/customerservice. Basic subscription rates: $59.94 per year within the United States; in all other countries, US$69. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051, Fort Erie, ON L2A 6C7. Printed in USA. Audited by the Alliance for Audited Media.