IS OCEAN ACIDIFICATION DRIVING SEA CREATURES CRAZY? PAGE 36

SCIENTIFIC AMERICAN

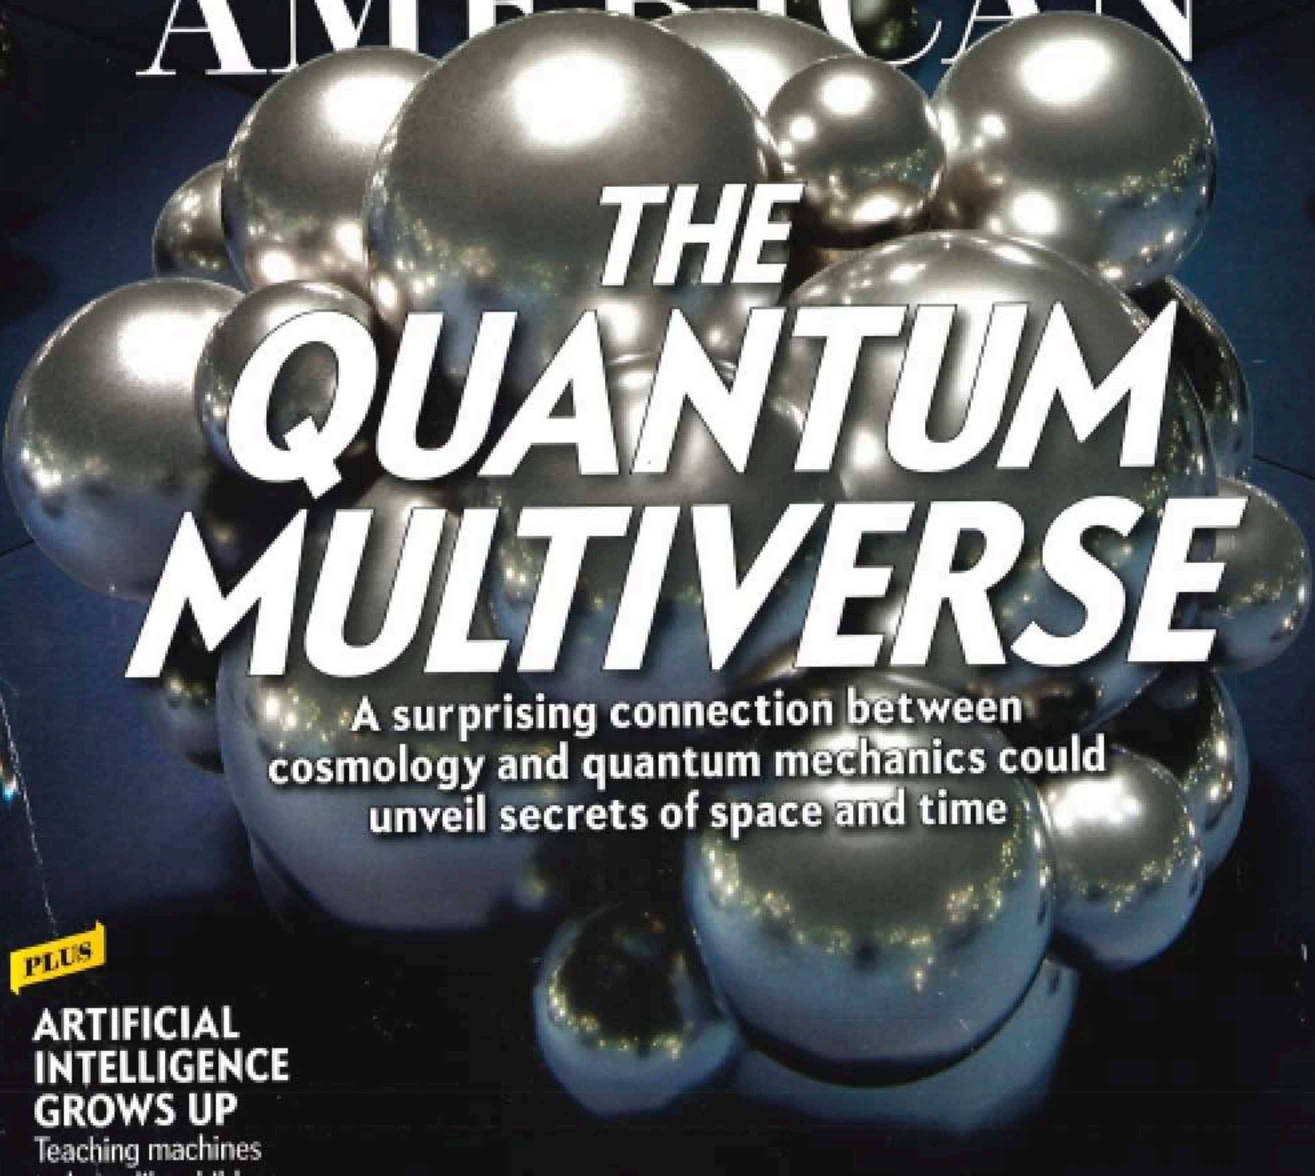

THE QUANTUM MULTIVERSE

A surprising connection between cosmology and quantum mechanics could unveil secrets of space and time

PLUS

ARTIFICIAL INTELLIGENCE GROWS UP

Teaching machines to learn like children PAGE 54

VANISHED VIKINGS

The mysterious fate of Greenland's Norse colonies PAGE 60

NEW HOPE FOR ALS

Gene therapies could hold promise PAGE 40

ScientificAmerican.com
U.K. £4.50 $5.99 U.S.
JUNE 2017

SCIENTIFIC AMERICAN

JUNE 2017

VOLUME 316, NUMBER 6

ON THE COVER

If our entire observable universe is only a bubble embedded in an infinitely larger multiverse, cosmologists may be in trouble. In a multiverse, all possible events occur an infinite number of times, stripping theories of predictive power. Ideas from quantum mechanics, however, offer fresh hope for more predictive cosmological theories. **Photograph by The Voorhes.**

SCIENTIFIC AMERICAN

8

10

70

ON THE WEB

Marching for Science
Watch *Scientific American*'s video coverage of the March for Science, held on April 22 in Washington, D.C.
Go to www.ScientificAmerican.com/jun2017/science-march

Scientific American (ISSN 0036-8733), Volume 316, Number 6, June 2017, published monthly by Scientific American, a division of Nature America, Inc., 1 New York Plaza, Suite 4500, New York, N.Y. 10004-1562. Periodicals postage paid at New York, N.Y., and at additional mailing offices. Canada Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40012504. Canadian BN No. 127387652RT; TVQ1218059275 TQ0001. Publication Mail Agreement #40012504. Return undeliverable mail to Scientific American, P.O. Box 819, Stn Main, Markham, ON L3P 8A2. **Individual Subscription rates:** 1 year $49.99 (USD), Canada $59.99 (USD), International $69.99 (USD). **Institutional Subscription rates:** Schools and Public Libraries: 1 year $84 (USD), Canada $89 (USD), International $96 (USD). Businesses and Colleges/Universities: 1 year $399 (USD), Canada $405 (USD), International $411 (USD). Postmaster: Send address changes to Scientific American, Box 3187, Harlan, Iowa 51537. **Reprints available:** write Reprint Department, Scientific American, 1 New York Plaza, Suite 4500, New York, N.Y. 10004-1562; fax: 646-563-7138; reprints@SciAm.com. **Subscription inquiries:** U.S. and Canada (800) 333-1199; other (515) 248-7684. **Send e-mail to** scacustserv@cdsfulfillment.com. Printed in U.S.A.

Science without Walls

Partnering across borders means faster discovery and a safer world

By the Editors

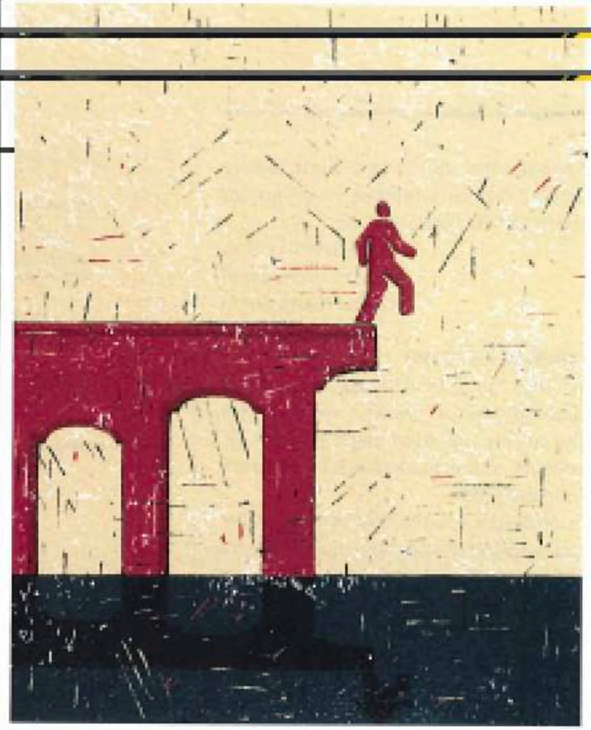

The U.S. appears to be plunging headlong into a new era of isolationism. The White House wants to pull out of international agreements, including the Paris climate deal and the North American Free Trade Agreement. It has issued executive orders trying to halt or slow the flow of refugees and immigrants to the nation.

This is bad for the U.S. and terrible for hundreds of thousands of desperate people across the planet. And it will strangle science. The choke hold will leave us more vulnerable to emerging, deadly viruses and will hamper efforts to explore space and control global threats such as climate change.

Research depends on ideas shared across political borders—including among countries in conflict. Even as the cold war was raging, hostility between the U.S. and the Soviet Union was put aside when American medical researcher Albert B. Sabin and his Soviet counterparts tested a live-virus, oral polio vaccine in the U.S.S.R. That successful trial provided the scientific proof needed for the vaccine's use around the world and ultimately helped to eradicate polio in most countries. During the International Polar Years of 1882–1883 and 1932–1933, nations also put aside their differences to study the Arctic and Antarctic.

Louis Pasteur once declared that "science knows no country, because knowledge belongs to humanity, and is the torch which illuminates the world." Nations have repeatedly seen the wisdom of his words.

The Soviets and Americans also worked together to further space exploration in the 1960s and 1970s—exchanging weather data from and launching new meteorologic satellites and jointly mapping the earth's geomagnetic field. Similarly, when the Soviet Union's Cosmos 936 mission launched in 1977, seven U.S. biological experiments were onboard. And in 2014, before the U.S.'s restoration of diplomatic relations with Cuba was in place, the American Association for the Advancement of Science and the Cuban Academy of Sciences pledged to work together to further research on drug resistance, cancer, emerging and infectious diseases, and the brain.

In recent years the U.S. has taken some crucial steps to strengthen our science diplomacy: In 2009 President Barack Obama spoke in Cairo about working with scientists in the Muslim world to develop novel sources of energy, create green jobs, digitize records, provide clean water and grow new crops. That speech led to the U.S. Science Envoy program, an outreach effort that selects top American scientists to promote the nation's commitment to science, technology and innovation as tools of diplomacy and economic growth abroad. One of the researchers in the program, vaccine scientist Peter Hotez, used his envoy position in the Middle East to create a vaccine research partnership between his American institute and a university in Saudi Arabia.

Yet the future of the envoy program under President Donald Trump remains unclear. Trump's travel bans have thrown researchers' plans into disarray—making foreign scientists and scholars question whether they should attempt to come to the U.S. for jobs or conferences and raising doubts about whether foreign scientists working here can risk visiting relatives in Muslim-majority countries, lest they be prevented from returning.

That is unfortunate because better science—and dialogue about science—benefits us all. Detecting and stopping emerging threats such as Zika or Ebola require partnering with countries around the globe. Understanding the extent of Zika damage and testing candidate vaccines among susceptible populations, for instance, will call for international cooperation.

For space exploration, we need Russia's assistance to ferry our astronauts to the International Space Station. To better map the stars and explore the unknown, we must partner with China because it has the world's largest radio telescope. To help limit the effects of climate change, we need all the big emitters, including the U.S., China and India, to take steps to address the issue and to work toward solutions that will help communities build resiliency.

Let's resist the urge to turn inward and isolate ourselves. Instead we must continue to forge strong ties worldwide, using science as a diplomatic wedge. We gain far more from these partnerships than we risk. Weakening them will hurt us all. ■

JOIN THE CONVERSATION ONLINE
Visit *Scientific American* on Facebook and Twitter
or send a letter to the editor: editors@sciam.com

Illustration by Augusto Zambonato

Robbert Dijkgraaf is director and Leon Levy Professor at the Institute for Advanced Study in Princeton, N.J. He is author of a companion essay in *The Usefulness of Useless Knowledge*, by Abraham Flexner (Princeton University Press, 2017).

Knowledge Is Infrastructure

Curiosity-driven science is just as vital as roads and bridges

By Robbert Dijkgraaf

When we think of infrastructure, we tend to think of the facilities and systems required for a country to function and thrive—roads, bridges, tunnels, airports and railways, as President Donald Trump specified in his February 28 speech to Congress.

Potholes and crumbling edifices clearly indicate that something needs fixing. But knowledge is infrastructure, too, and right now it needs urgent attention. Science and technology are the basis of the modern economy and key to solving many serious environmental, social and security challenges. Basic research, driven by curiosity, freedom and imagination, provides the groundwork for all applied research and technology. And just as we have to break the endless cycle of Band-Aid fixes to roads and rails, long-term investments in knowledge are vital.

Curiosity-driven basic research has brought truly revolutionary transformations, such as the rapid growth of computer-based intelligence and the discovery of the genetic basis of life. Albert Einstein's century-old theory of relativity is used every day in our GPS devices. Perhaps the best U.S. government investment ever was the $4.5-million grant from the National Science Foundation that led to the Google search algorithm—an investment that has multiplied by more than 100,000 times.

Basic research not only radically alters our deep understanding of the world, it also leads to new tools and techniques that spread throughout society, such as the World Wide Web, originally developed for particle physicists to foster scientific collaboration. It trains the sharpest minds on the toughest challenges, and its products are widely used by industry and society. No one can exclusively capture its rewards—it is a truly public good.

The path from exploratory basic research to practical applications is not one-directional and linear, but rather it is complex and cyclic. The resulting technologies enable even more fundamental discoveries, such as quantum mechanics, which has led to computer chips and other inventions that are responsible for a significant portion of the U.S. gross domestic product (GDP).

To tap into the full potential of human intellect and imagination, we need to balance short-term expectations with long-term investment. Just as a financial expert would never recommend forgoing a retirement fund to enrich an already sufficient checking account, we need to advocate for a balanced portfolio of short- and long-term research initiatives. But driven by decreasing funding, against a background of economic uncertainty, global political turmoil and ever shortening time cycles, research is becoming dangerously skewed toward short-term goals that may address

current problems but miss out on huge advances in the long term.

It is a worrisome trend that over the past decades both public and private support for basic research have declined as a percentage of GDP. The postwar decades saw an unprecedented worldwide growth of science, including the creation of funding councils such as the National Science Foundation and massive investments in research infrastructure. Recent years have seen a marked retrenchment. Steadily declining public funding is insufficient to keep up with the expanding role of the scientific enterprise in a modern knowledge-based society. The U.S. federal R&D budget, measured as a fraction of GDP, has dropped from a high of 1.92 percent in 1964, at the height of the cold war and the space race, to less than 0.8 percent today. And the budget for the National Institutes of Health has fallen since 2003.

Governments are increasingly directing research funding to tackle important societal challenges, such as transitioning to clean, sustainable energy, battling climate change and preventing worldwide epidemics, all within flat or decreasing budgets. As a consequence, basic research and its budget are given short shrift.

It is human to focus on necessities in times of stress. But investing in basic research, just like saving for retirement, is a prerequisite for ensuring welfare, innovation and societal progress. Long-term investments in basic research are crucial and lead to an even higher goal: the global benefits of embracing the scientific culture of accuracy, truth seeking, critical questioning and dialogue, healthy skepticism, respect for facts and uncertainties, and wonder at the richness of nature and the human spirit. ■

JOIN THE CONVERSATION ONLINE
Visit *Scientific American* on Facebook and Twitter or send a letter to the editor: editors@sciam.com

Illustration by Bernard Lee

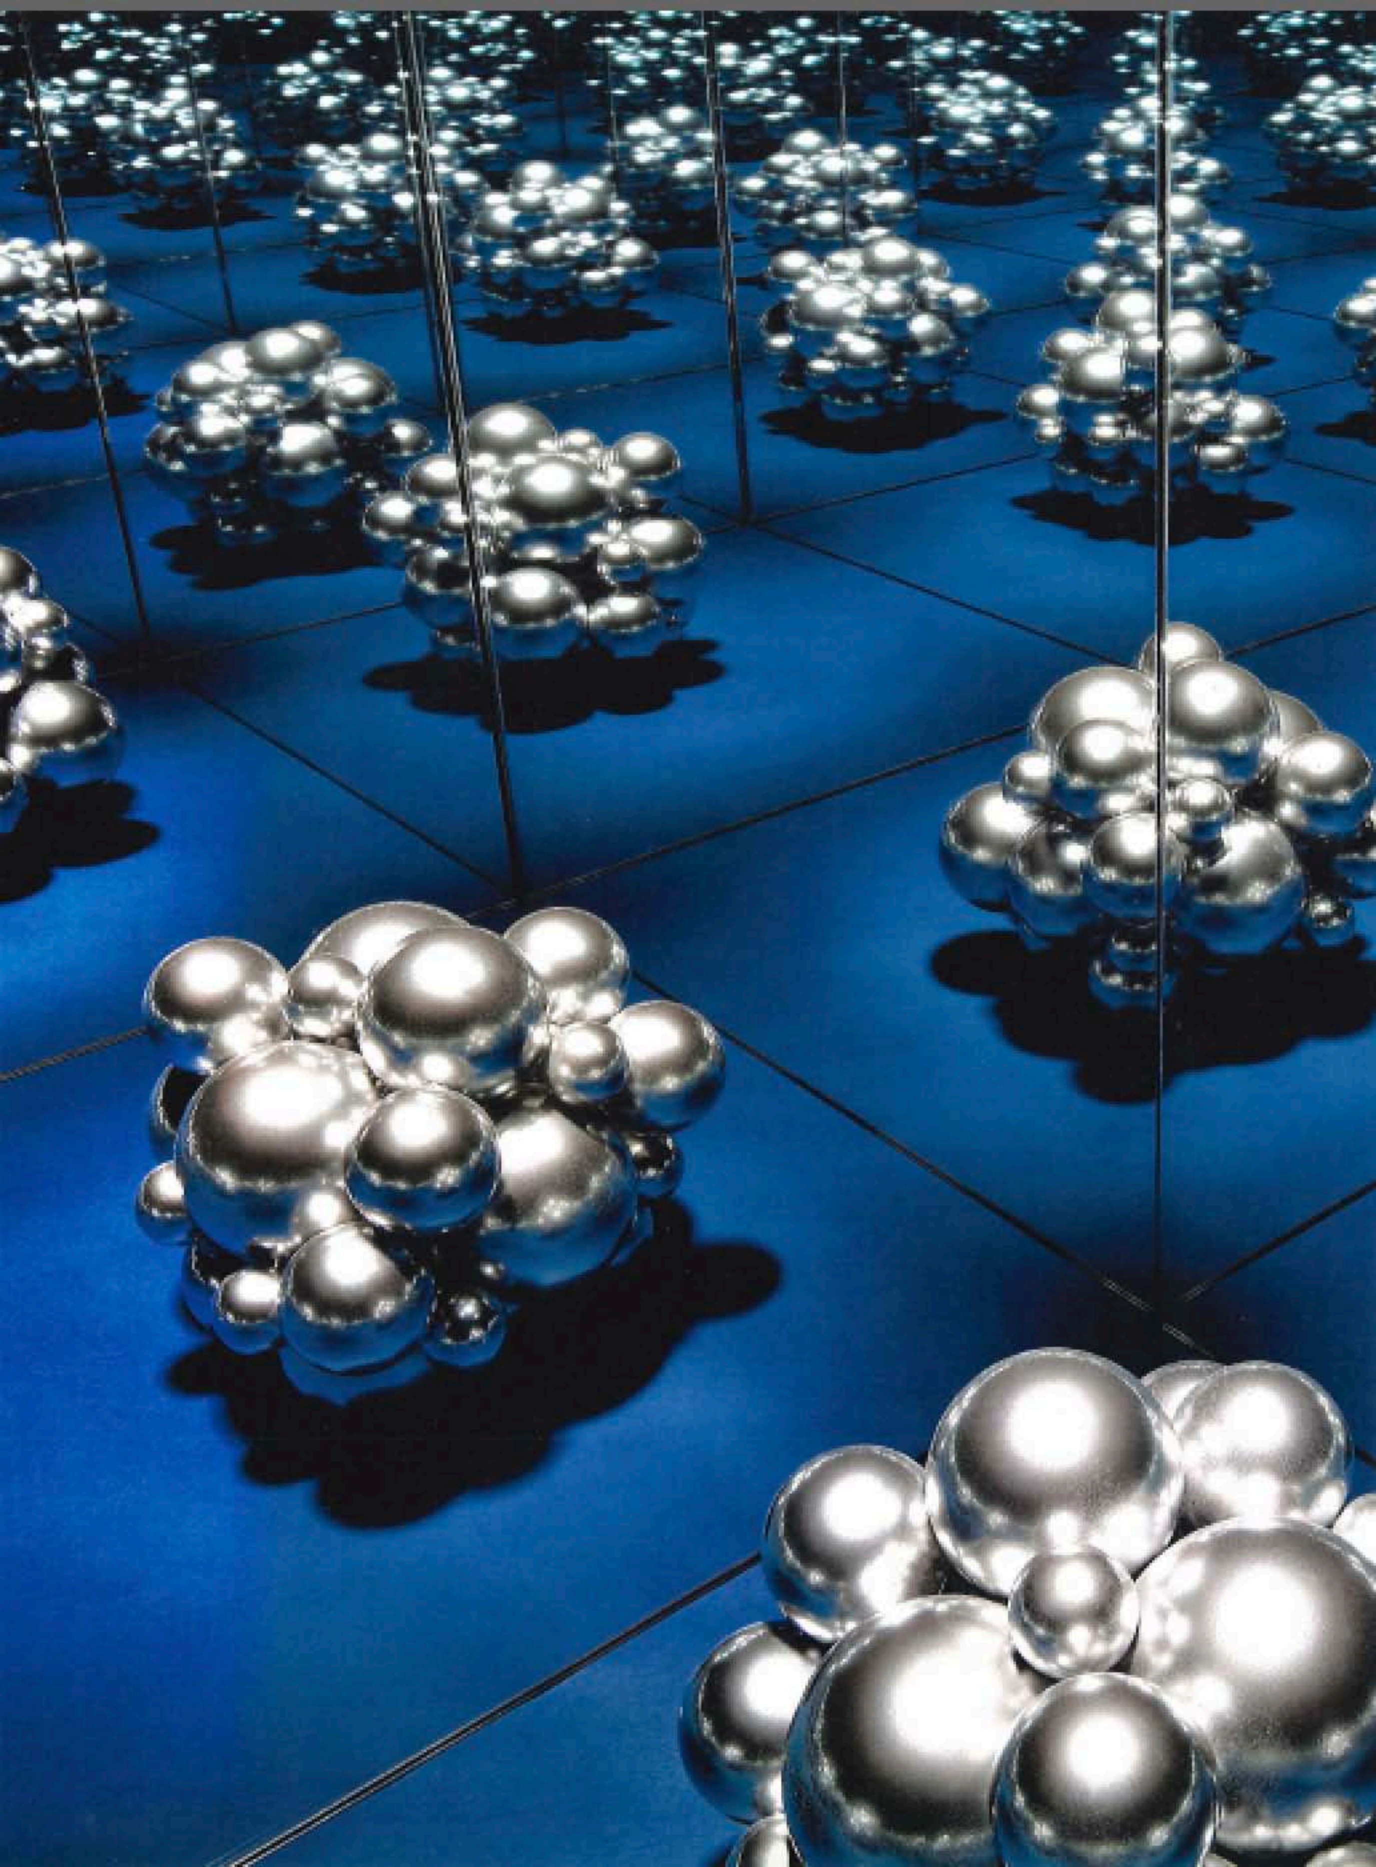

THE QUANTUM MULTIVERSE

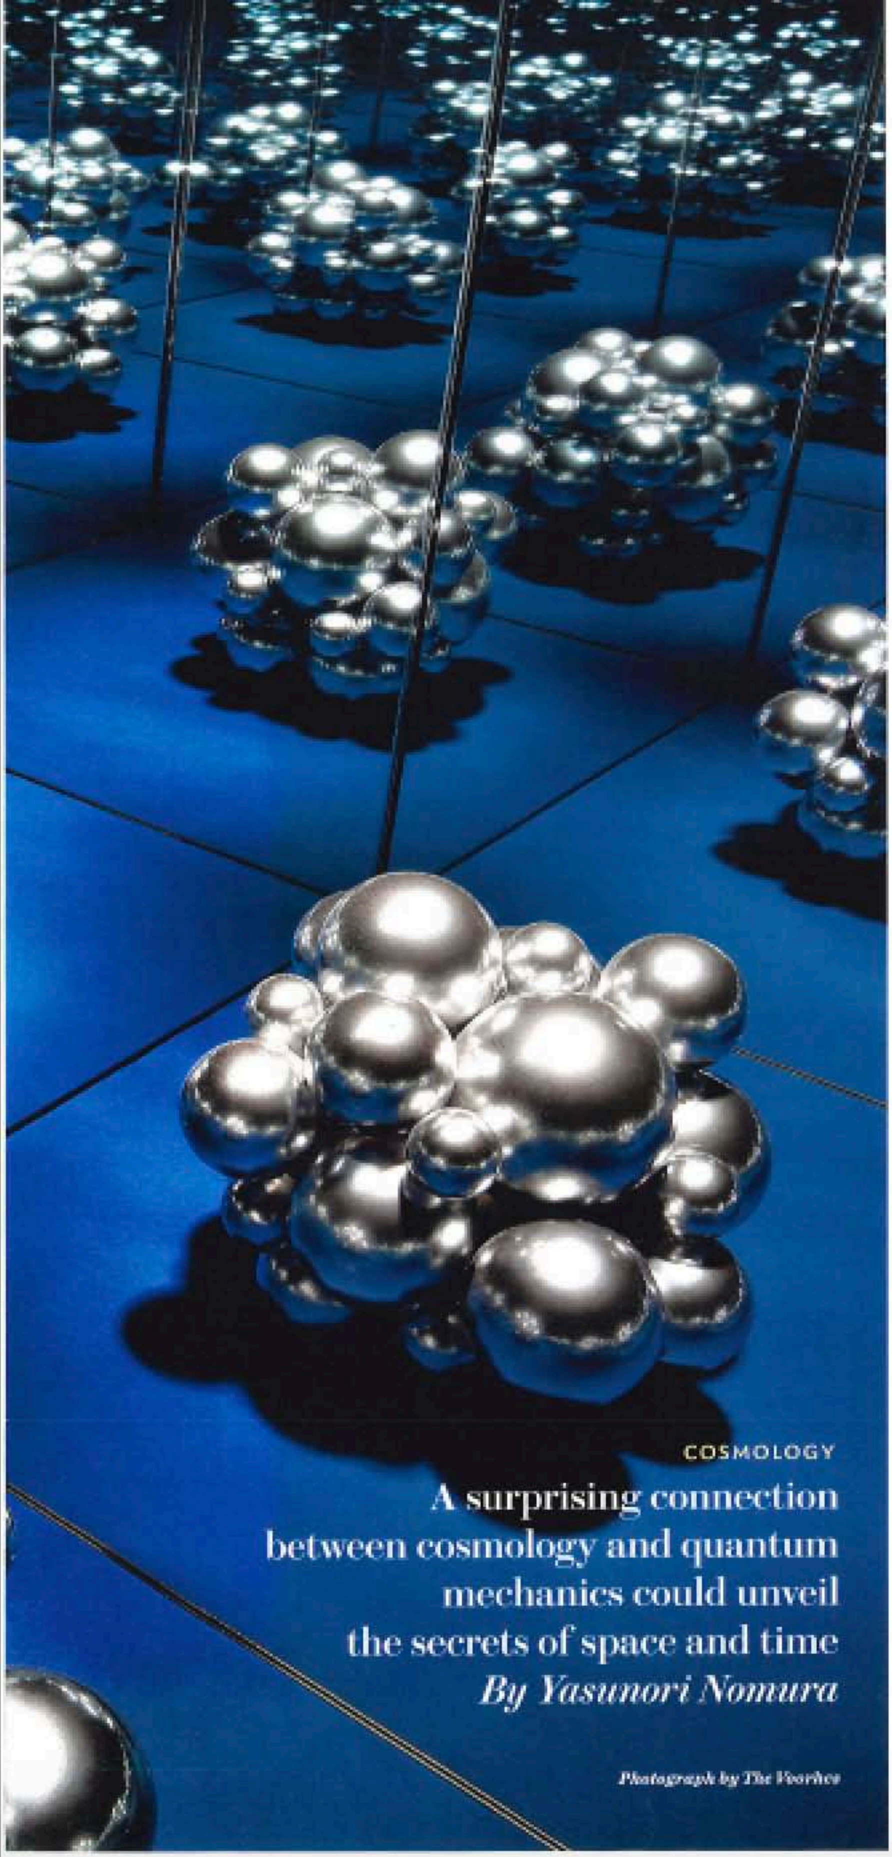

Photograph by The Voorhes

COSMOLOGY

A surprising connection between cosmology and quantum mechanics could unveil the secrets of space and time

By Yasunori Nomura

Yasunori Nomura is a professor of physics and director of the Berkeley Center for Theoretical Physics at the University of California, Berkeley. He is also a senior faculty scientist at Lawrence Berkeley National Laboratory and a principal investigator at the University of Tokyo's Kavli Institute for the Physics and Mathematics of the Universe.

MANY COSMOLOGISTS NOW ACCEPT THE EXTRAORDINARY IDEA THAT what seems to be the entire universe may actually be only a tiny part of a much larger structure called the multiverse. In this picture, multiple universes exist, and the rules we once assumed were basic laws of nature take different forms in each; for example, the types and properties of elementary particles may differ from one universe to another.

The multiverse idea emerges from a theory that suggests the very early cosmos expanded exponentially. During this period of "inflation," some regions would have halted their rapid expansion sooner than others, forming what are called bubble universes, much like bubbles in boiling water. Our universe would be just one of these bubbles, and beyond it would lie infinitely more.

The idea that our entire universe is only a part of a much larger structure is, by itself, not as outlandish as it sounds. Throughout history scientists have learned many times over that the visible world is far from all there is. Yet the multiverse notion, with its unlimited number of bubble universes, does present a major theoretical problem: it seems to erase the ability of the theory to make predictions—a central requirement of any useful theory. In the words of Alan Guth of the Massachusetts Institute of Technology, one of the creators of inflation theory, "in an eternally inflating universe, anything that can happen will happen; in fact, it will happen an infinite number of times."

In a single universe where events occur a finite number of times, scientists can calculate the relative probability of one event occurring versus another by comparing the number of times these events happen. Yet in a multiverse where everything happens an infinite number of times, such counting is not pos-

physicists. Some researchers, including me, have now realized that quantum theory—which, in contrast to the multiverse notion, is concerned with the very smallest particles in existence—may, ironically, point the way to a solution. Specifically, the cosmological picture of the eternally inflating multiverse may be mathematically equivalent to the "many worlds" interpretation of quantum mechanics, which attempts to explain how particles can seem to be in many places at once. As we will see, such a connection between the theories not only solves the prediction problem, it may also reveal surprising truths about space and time.

QUANTUM MANY WORLDS

I CAME TO THE IDEA of a correspondence between the two theories after I revisited the tenets of the many-worlds interpretation of quantum mechanics. This concept arose to make sense of some of the stranger aspects of quantum physics. In the quantum world—a nonintuitive place—cause and effect work differently than they do in the macro world, and the outcome of any process is always probabilistic. Whereas in our macroscopic experience, we can predict where a ball will land when it is thrown based on its starting point, speed and other factors, if that ball were a quantum particle, we could only ever say it has a certain

Inflation Meets Many Worlds

The theory of inflation suggests that our universe is one of infinitely many that formed when the very early cosmos expanded exponentially. This picture of a multiverse, however, seems to destroy the theory's ability to make predictions because anything that can happen in an infinite multiverse will happen infinitely many times. The problem is solved, however, if the inflationary multiverse is equivalent to the "many worlds" interpretation of quantum mechanics, which posits that all these infinite universes coexist not in a single real space but in "probability space."

MANY WORLDS

Quantum mechanics says that a particle, rather than being hidden under either cup A or cup B, actually exists under both cups with a certain probability (*denoted by yellow wave*) of being found in any given place. Only when an observer turns over the cups to check does the particle "choose" to be in one of the two possible locations. The many-worlds interpretation suggests that every time an observer performs such a measurement, two new universes branch off—one where the particle ended up being under cup A and one where the particle resided under cup B.

INFLATIONARY MULTIVERSE

This theory holds that during inflation certain regions would have slowed their rapid expansion before others, forming bubbles that became universes unto themselves. As time went on, more and more patches slowed to form new bubbles within the larger inflating space, which went on expanding eternally. Our universe is just one of these bubbles.

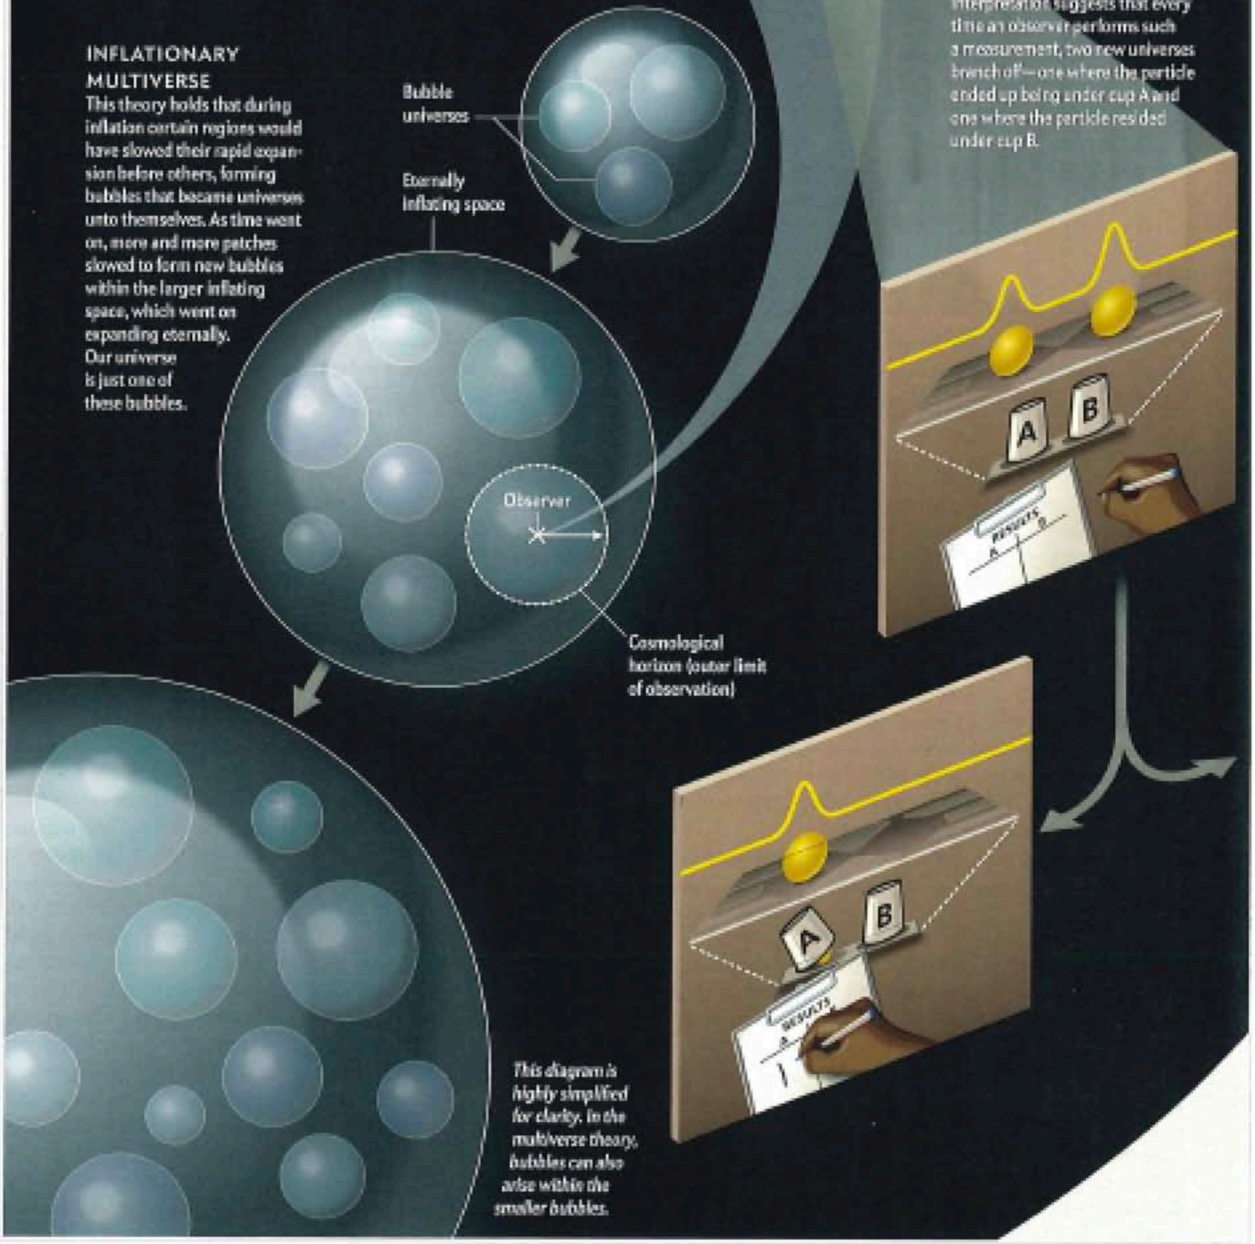

This diagram is highly simplified for clarity. In the multiverse theory, bubbles can also arise within the smaller bubbles.

Illustration by Jen Christiansen

THE MESSY TRUTH ABOUT WEIGHT LOSS

HEALTH

Two decades of research confirm that weight loss is about burning more calories than you consume—but what you eat is more important than how much you exercise

By Susan B. Roberts and Sai Krupa Das

Photograph by Dan Saelinger

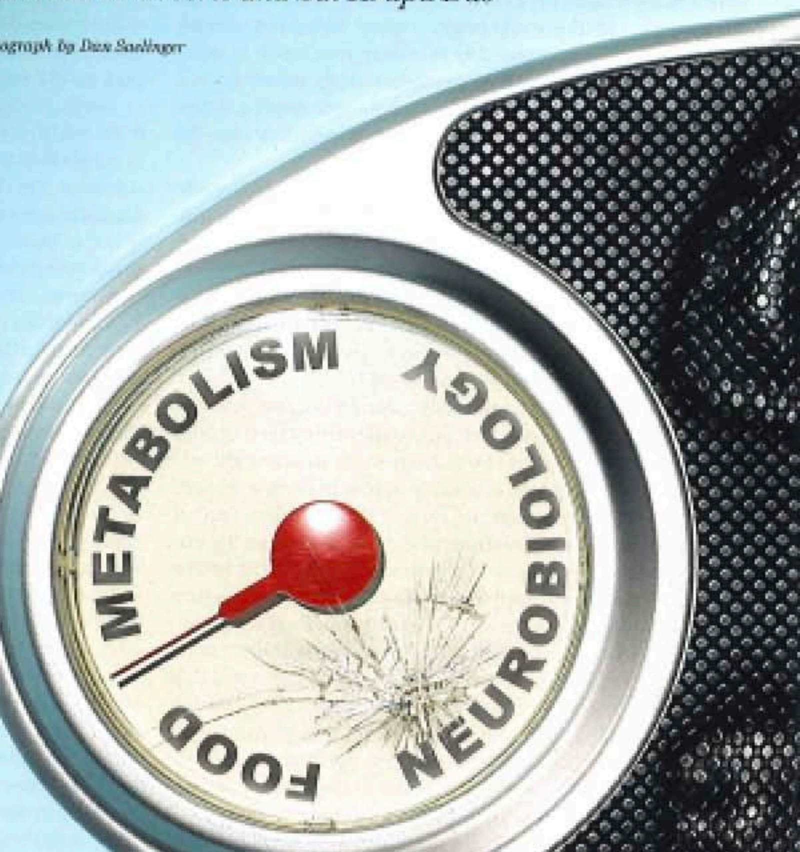

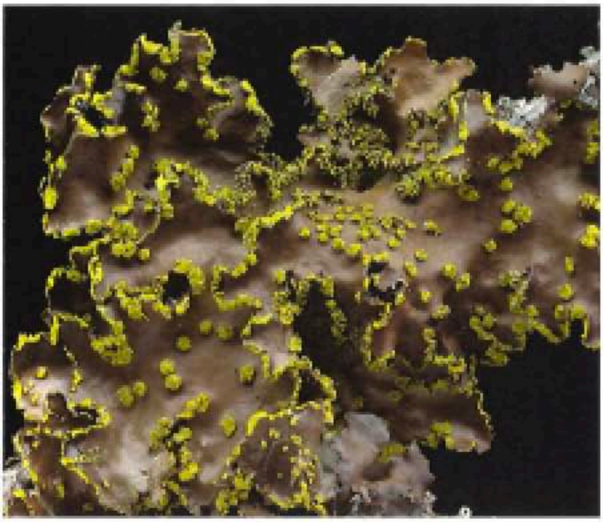

BIOLOGY

The Meaning of Lichen

How a self-taught naturalist unearthed hidden symbioses in the wilds of British Columbia—and helped to overturn 150 years of accepted scientific wisdom

By Erica Gies

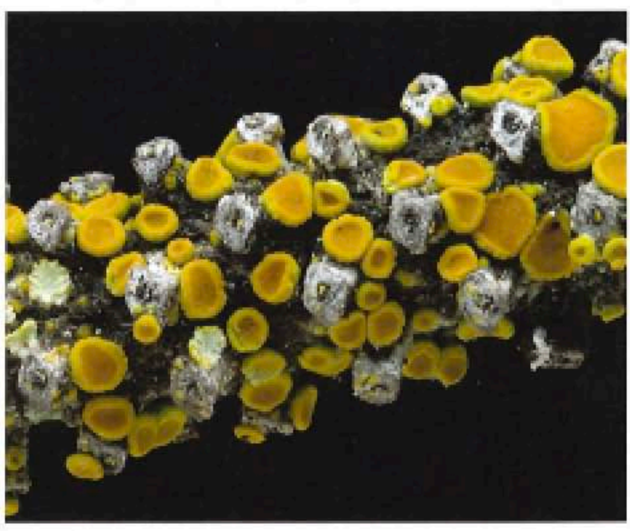

MACHINE LEARNING

MAKING AI MORE HUMAN

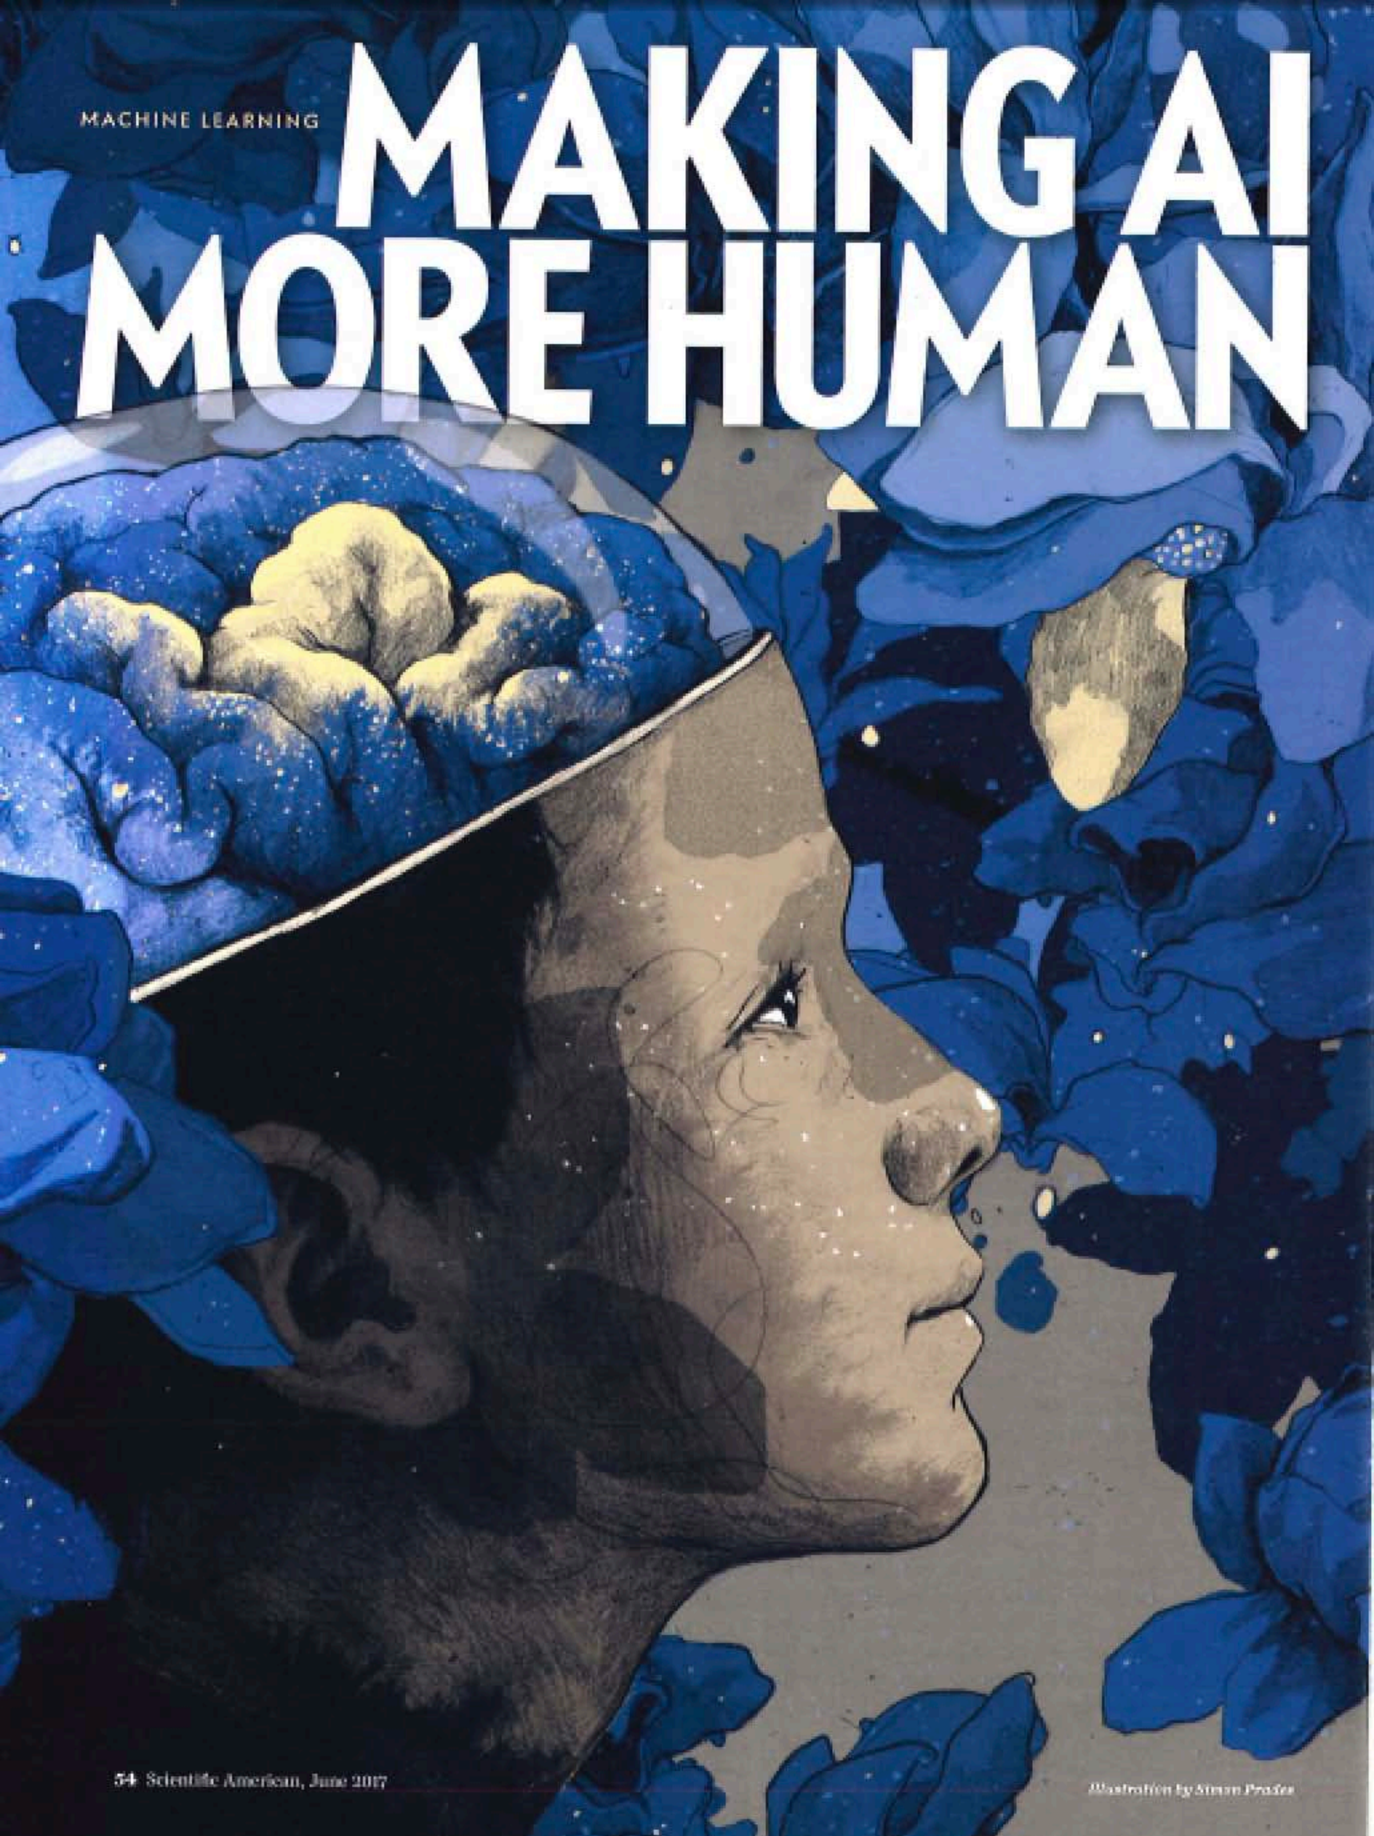

Illustration by Simon Prades

Artificial intelligence has staged a revival by starting to incorporate what we know about how children learn

By Alison Gopnik

Alison Gopnik is a professor of psychology and an affiliate professor of philosophy at the University of California, Berkeley. Her research focuses on how young children learn about the world around them.

IF YOU SPEND MUCH TIME WITH CHILDREN, YOU'RE BOUND TO WONDER HOW young human beings can possibly learn so much so quickly. Philosophers, going all the way back to Plato, have wondered, too, but they've never found a satisfying answer. My five-year-old grandson, Augie, has learned about plants, animals and clocks, not to mention dinosaurs and spaceships. He also can figure out what other people want and how they think and feel. He can use that knowledge to classify what he sees and hears and make new predictions. He recently proclaimed, for example, that the newly discovered species of titanosaur on display at the American Museum of Natural History in New York City is a plant eater, so that means it really isn't that scary.

Yet all that reaches Augie from his environment is a stream of photons hitting his retina and disturbances of air contacting his eardrums. The neural computer that sits behind his blue eyes manages somehow to start with that limited information from his senses and to end up making predictions about plant-eating titanosaurs. One lingering question is whether electronic computers can do the same.

During the past 15 years or so computer scientists and psychologists have been trying to find an answer. Children acquire a great deal of knowledge with little input from teachers or parents. Despite enormous strides in machine intelligence, even the most powerful computers still cannot learn as well as a five-year-old does.

Figuring out how the child brain actually functions—and then creating a digital version that will work as effectively—will challenge computer scientists for decades to come. But in the meantime, they are beginning to develop artificial intelligence that incorporates some of what we know about how humans learn.

THIS WAY UP

AFTER THE FIRST BURST of enthusiasm in the 1950s and 1960s, the quest for AI languished for decades. In the past few years, though, there have been striking advances, especially in the field of machine learning, and AI has become one of the hottest developments in technology. Many utopian or apocalyptic predictions have emerged about what those advances mean. They have, quite literally, been taken to presage either immortality or the end of the world, and a lot has been written about both these possibilities.

I suspect that developments in AI lead to such strong feelings because of our deep-seated fear of the almost human. The idea that creatures might bridge the gap between the human and the artificial has always been deeply disturbing, from the medieval golem to Frankenstein's monster to Ava, the sexy robot fatale in the movie *Ex Machina*.

But do computers really learn as well as humans? How much of the heated rhetoric points to revolutionary change, and how much is just hype? The details of how computers learn to recognize, say, a cat, a spoken word or a Japanese character can be hard to follow. But on closer inspection, the basic ideas behind machine learning are not as baffling as they first seem.

One approach tries to solve the problem by starting with the stream of photons and air vibrations that Augie, and all of us, receives—and that reaches the computer as pixels of a digital image and sound samples of an audio recording. It then tries to extract a series of patterns in the digital data that can detect and identify whole objects in the surrounding world. This so-called bottom-up approach has roots in the ideas of philosophers such as David Hume and John Stuart Mill and psychologists such as Ivan Pavlov and B. F. Skinner, among others.

In the 1980s scientists figured out a compelling and inge-

IN BRIEF

How do young children know what they know? That question has long preoccupied philosophers and psychologists—and now computer scientists.

Specialists in artificial intelligence are studying the mental reasoning powers of preschoolers to develop ways to teach machines about the world.

Two rival machine-learning strategies—both halting attempts to mimic what children do naturally—have begun to transform AI as a discipline.

Two Paths to AI's Resurgence

Problems the average five-year-old solves readily can stump even the most powerful computers. AI has made a spirited comeback in recent years by teaching computers to learn about the world somewhat like a child does. The machine recognizes the letter "A" either from raw sensory information—a bottom-up approach—or by making a guess from preexisting knowledge—a top-down approach.

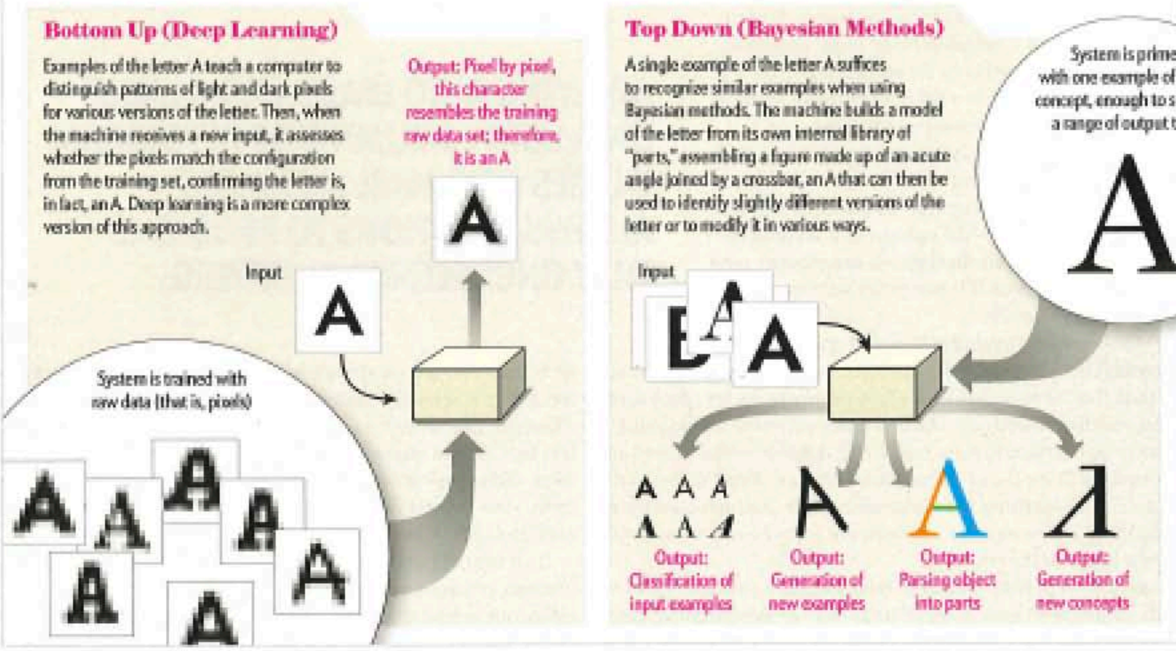

SOURCE: "HUMAN-LEVEL CONCEPT LEARNING THROUGH PROBABILISTIC PROGRAM INDUCTION," BY BRENDEN M. LAKE, RUSLAN SALAKHUTDINOV AND JOSHUA B. TENENBAUM, IN SCIENCE, VOL. 350; DECEMBER 11, 2015 (Bayesian example)

nious way to apply bottom-up methods to let computers hunt for meaningful patterns in data. "Connectionist," or "neural network," systems take inspiration from the way that neurons convert light patterns at your retina into representations of the world around you. A neural network does something similar. It uses interconnected processing elements, akin to biological cells, to transform pixels at one layer of the network into increasingly abstract representations—a nose or an entire face—as data are crunched at progressively higher layers.

Neural-network ideas have gone through a recent revival because of new techniques called deep learning—technology now being commercialized by Google, Facebook and other tech giants. The ever increasing power of computers—the exponential increase in computing capability that is captured by what is known as Moore's law—also has a part in the new success of these systems. So does the development of enormously large data sets. With better processing capabilities and more data to crunch, connectionist systems can learn far more effectively than we might have once thought.

Over the years the AI community has seesawed between favoring these kinds of bottom-up solutions to machine learning and alternative top-down approaches. Top-down approaches leverage what a system already knows to help it learn something new. Plato, as well as so-called rationalist philosophers such as René Descartes, believed in a top-down approach to learning—and it played a big role in early AI. In the 2000s such methods also experienced their own rebirth in the form of probabilistic, or Bayesian, modeling.

Like scientists, top-down systems start out by formulating abstract and wide-ranging hypotheses about the world. The systems then make predictions about what the data should look like if those hypotheses are correct. Also like scientists, the systems then revise their hypotheses, depending on the outcome of those predictions.

NIGERIA, VIAGRA AND SPAM

BOTTOM-UP METHODS are perhaps the most readily understood, so let's consider them first. Imagine that you are trying to get your computer to separate important messages from the spam that arrives in your in-box. You might notice that spam tends to have certain distinguishing characteristics: a long list of recipient addressees, an originating address in Nigeria or Bulgaria, references to $1-million prizes or perhaps mention of Viagra. But perfectly useful messages might look the same. You don't want to miss the announcement that you have earned a promotion or an academic award.

If you compare enough examples of spam against other types of e-mails, you might notice that only the spam tends to have qualities that combine in certain telltale ways—Nigeria, for instance, plus a promise of a $1-million prize together spell

Illustration by Jen Christiansen

trouble. In fact, there might be some quite subtle higher-level patterns that discriminate between the spam messages and the useful ones—misspellings and IP addresses that are not at all obvious, for example. If you could detect them, you could accurately filter out the spam—without fear of missing a notice that your Viagra has shipped.

Bottom-up machine learning can ferret out the relevant clues to solve this kind of task. To do this, a neural network must go through its own learning process. It evaluates millions of examples from huge databases, each labeled as spam or as an authentic e-mail. The computer then extracts a set of identifying features that separate spam from everything else.

In a similar way, the network might inspect Internet images labeled "cat," house," "stegosaurus," and so on. By extracting the common features in each set of images—the pattern that distinguishes all the cats from all the dogs—it can identify new images of a cat, even if it has never seen those particular images before.

One bottom-up method, called unsupervised learning, is still in its relative infancy, but it can detect patterns in data that have no labels at all. It simply looks for clusters of features that identify an object—noses and eyes, for example, always go together to form a face and differ from the trees and mountains in the background. Identifying an object in these advanced deep-learning networks takes place through a division of labor in which recognition tasks are apportioned among different layers of the network.

An article in *Nature* in 2015 demonstrated just how far bottom-up methods have come. Researchers at DeepMind, a company owned by Google, used a combination of two different bottom-up techniques—deep learning and reinforcement learning—in a way that enabled a computer to master Atari 2600 video games. The computer began knowing nothing about how the games worked. At first, it made random guesses about the best moves while receiving constant feedback about its performance. Deep learning helped the system identify the features on the screen, and reinforcement learning rewarded it for a high score. The computer achieved a high proficiency level with several games; in some cases, it performed better than expert human players. That said, it also completely bombed on other games that are just as easy for humans to master.

The ability to apply AI to learn from large data sets—millions of Instagram images, e-mail messages or voice recordings—allows solutions to problems that once seemed daunting, such as image and speech recognition. Even so, it is worth remembering that my grandson has no trouble at all recognizing an animal or responding to a spoken query even with much more limited data and training. Problems that are easy for a human five-year-old are still extremely perplexing to computers and much harder than learning to play chess.

Computers that learn to recognize a whiskered, furry face often need millions of examples to categorize objects that we can classify with just a few. After extensive training, the computer might be able to identify an image of a cat that it has never seen before. But it does so in ways that are quite different from human generalizations. Because the computer software reasons differently, slipups occur. Some cat images will not be labeled as cats. And the computer may incorrectly say an image is a cat, although it is actually just a random blur, one that would never fool a human observer.

APPLYING AI TO LEARN FROM LARGE DATA SETS—MILLIONS OF INSTAGRAM IMAGES OR E-MAIL MESSAGES—ALLOWS SOLUTIONS TO PROBLEMS THAT ONCE SEEMED DAUNTING.

ALL THE WAY DOWN

THE OTHER APPROACH to machine learning that has transformed AI in recent years works in the opposite direction, from the top down. It assumes that we can get abstract knowledge from concrete data because we already know a lot and especially because the brain is already capable of understanding basic abstract concepts. Like scientists, we can use those concepts to formulate hypotheses about the world and make predictions about what data (events) should look like if those hypotheses are right—the reverse of trying to extract patterns from the raw data themselves, as in bottom-up AI.

This idea can best be illustrated by revisiting the spam plague through considering a real case in which I was involved. I received an e-mail from the editor of a journal with a strange name, referring specifically to one of my papers and proposing that I write an article for the publication. No Nigeria, no Viagra, no million dollars—the e-mail had none of the common indications of a spam message. But by using what I already knew and thinking in an abstract way about the process that produces spam, I could figure out that this e-mail was suspicious.

To start, I knew that spammers try to extract money from people by appealing to human greed—and academics can be as greedy to publish as ordinary folks are for $1-million prizes or better sexual performance. I also knew that legitimate "open access" journals have started covering their costs by charging authors instead of subscribers. Also, my work has nothing to do with the journal title. Putting all that together, I produced a plausible hypothesis that the e-mail was trying to sucker academics into paying to "publish" an article in a fake journal. I could draw this conclusion from just one example, and I could go on to test my hypothesis further by checking the editor's bona fides through a search-engine query.

A computer scientist would call my reasoning process a "generative model," one that is able to represent abstract concepts, such as greed and deception. This same model can also describe the process that is used to come up with a hypothesis—the reasoning that led to the conclusion that the message might be an e-mail scam. The model lets me explain how this form of spam works, but it also lets me imagine other kinds of spam or even a type that differs from any I have seen or heard about before. When I receive the e-mail from the journal, the model lets me work backward—tracing step by step why it must be spam.

SCIENTIFIC AMERICAN ONLINE Read more about how researchers use Bayesian methods at ScientificAmerican.com/jun2017/gopnik