10 BREAKTHROUGH TECHNOLOGIES

MIT Technology Review

SPECIAL GUEST CURATOR **BILL GATES** PICKS THIS YEAR'S LIST

New Wave Nuclear Power

Carbon Dioxide Catcher

Smooth-Talking AI Assistants

An ECG on Your Wrist

Predicting Preemies

Gut Probe in a Pill

Custom Cancer Vaccines

The Cow-Free Burger

Sanitation without Sewers

Robot Dexterity

Vol.122 No.2 $9.99 USD
Mar/Apr 2019 $10.99 CAD

There's lots more on this guy's mind.
See page 56.


Connecting vehicles
Protecting pedestrians
Making cities smarter
Reducing accidents
Maximizing efficiency
battelle.org/tech


SUNY State College of Optometry
BATTELLE
It can be done

In person, Bill Gates has something of both the sage and the child about him. His encyclopedic knowledge is legendary, and the quizzical furrow of his brow when you formulate a question unclearly hints at an impatience with lesser intellects. But get him talking on a subject that interests him—which is just about any subject under the sun—and you sense that he has never really stopped being the nerdy teenager in awe at the richness and complexity of the world he is exploring.

When a chance conversation led to the proposal that he choose MIT Technology Review's annual list of 10 breakthrough technologies, we were thrilled, but also, in hindsight, a little complacent. We've been compiling these lists since 2001, and we thought that if we offered Bill a shortlist of 20 to choose from, he would pick 10 and be done with it.

He rejected almost of all of them.

This list, then, is very much Bill's own, and as he explains in his introduction (page 8) and my interview with him (page 56), it represents a singularly Gatesian belief: that for all the ills remaining in the world, human welfare has made so much progress that we are now moving through a slow technological tipping point. If in the past most breakthroughs were about making life longer, in the future most will be about making it more agreeable. It's a bold and optimistic view—Bill is nothing if not an optimist—and whether or not you share it, it provides an interesting lens through which to look at the big technological trends of today.

Gideon Lichfield is editor in chief of MIT Technology Review.

IAN ALLEN

Bill's list focuses on three broad areas: climate change, health care, and AI. Not surprisingly, many of the items are related either to his charitable foundation's work or to his own investments. We've disclosed those relationships, but whereas for a journalist they'd constitute a conflict of interest, in Bill's case they reflect his own beliefs about which technologies will do the most good for humanity, which is precisely why we asked his opinion. It would be strange if he weren't investing in some of them.

To complement Bill's list we've compiled some of our own: 10 grand challenges that technology has yet to solve (page 18), 10 low-tech solutions that have had a big impact (page 22), and 10 of this century's biggest technology failures (page 88)—a list that, it turns out, was harder to agree on than we thought.

As in past years, we've featured some of the 10 breakthrough technologies in greater depth. The rest of the articles in the issue all look, in one way or another, at how innovation happens. Dayna Evans (page 78) shows the barriers that certain groups of entrepreneurs still face in her profile of a women's-health startup. David Rotman (page 58) examines how AI could revitalize industries like pharma and materials, where new breakthroughs are getting increasingly expensive. Brian Bergstein (page 82) looks at how non-tech companies like perfume makers are starting to adopt AI to help them innovate, and why it's usually much harder than they expect. Kate Chandler, who researches drone use in Africa, talks (page 76) about the pitfalls of importing a technology solution to the developing world without understanding the local context. David Silver, creator of AlphaGo and its successors, muses (page 66) on what it means for an AI to exhibit creativity, while Harvard philosopher Sean Dorrance Kelly (page 68) argues that machine creativity can never substitute for the human variety.

As always, we hope you find the list thought-provoking, and I'm interested in your thoughts on what made the cut (or what didn't). Write to me at gideon.lichfield@technologyreview.com and let me know.


©2019 America's Biopharmaceutical Companies.
The day we control severe asthma
is the day I breathe easier.
René / Asthma Researcher
Gil / Asthma Patient
As the world leader in medical innovation, America's researchers are developing 54 new medicines to treat or prevent asthma. Because for over 24 million Americans living with asthma, the only way to stop the next attack is to attack first.
Innovation.org
GOBOLDLY
America's Biopharmaceutical Companies

Contents

10 BREAKTHROUGH TECHNOLOGIES

THE 2019 list

COVER PHOTOGRAPH BY IAN ALLEN; LETTERING BY CHRIS PIASCIK

Editorial

Editor in chief
Gideon Lichfield

Deputy editor
Michael Reilly

Editor at large
David Rotman

News editor
Niall Firth

Managing editor
Timothy Maher

Commissioning editors
Bobbie Johnson,
Konstantin Kakaes

San Francisco bureau chief
Martin Giles

Senior editor, MIT News
Alice Dragoon

Senior editor, AI and robotics
Will Knight

Senior editor, biomedicine
Antonio Regalado

Senior editor, energy
James Temple

Senior editor, infrastructure
Elizabeth Woyke

Reporters
Karen Hao (AI),
Charlotte Jee (news),
Mike Orcutt (blockchain),
Erin Winick (space)

Copy chief
Linda Lowenthal

Social media associate
Benji Rosen

Senior production director
James LaBelle

Research assistant
Tate Ryan-Mosley

Administrative assistant
Andrea Siegel

Design

Chief creative officer
Eric Mongeon

Art director
Emily Luong

Marketing and events designer
Kyle Thomas Hemingway

Assistant art director
Emily Caulfield

Corporate

Chief executive officer and publisher
Elizabeth Bramson-Boudreau

Human resources manager
James Wall

Assistant to the CEO
Katie McLean

Manager of information technology
Colby Wheeler

Office manager
Linda Cardinal

Product development

Chief digital officer
Cy Caine

Director of product
Vanessa DeCollibus

Project manager
Allison Chase

UI/UX designer
Jon Akland

Engineers
Shaun Calhoun,
Molly Frey,
Jason Lewicki

Licensing and communities

Vice president, licensing and communities
Antoinette Matthews

Client services manager
Ted Hu

Events

Senior vice president,
events and strategic partnerships
Amy Lammers

Director of events programming
Laura Janes Wilson

Senior events manager
Nicole Silva

Content and program developer, events
Kelsie Pallanck

Event content producer,
custom and international
Marcy Rizzo

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
Tom Russell

Director of audience development
Rosemary Kelly

Director of digital marketing
Josh Getman

Assistant consumer marketing manager
Sanjeet Chowdhury

Assistant product marketing manager
Amanda Saeli

Advertising sales

Vice president, sales and brand partnerships
Marii Sebahar
marii@technologyreview.com
415-416-9140

Senior 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

Digital sales strategy 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

Senior editor
Mindy Blodgett

Senior project manager
Anna Raborn

Content manager
Jason Sparapani

Director of consulting, Asia
Claire Beatty

Director of business development, Asia
Marcus Ulvne

Customer service and subscription inquiries

National
800-877-5230

International
903-636-1115

E-mail
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.

Portrait by **IAN ALLEN**

How we'll invent the future

BREAKTHROUGH 20 10 19 TECHNOLOGIES

BY **Bill Gates**

I was honored when MIT Technology Review invited me to be the first guest curator of its 10 Breakthrough Technologies. Narrowing down the list was difficult. I wanted to choose things that not only will create headlines in 2019 but captured this moment in technological history—which got me thinking about how innovation has evolved over time.

My mind went to—of all things—the plow. Plows are an excellent embodiment of the history of innovation. Humans have been using them since 4000 BCE, when Mesopotamian farmers aerated soil with sharpened sticks. We've been slowly tinkering with and improving them ever since, and today's plows are technological marvels.

By **BILL GATES**

But what exactly is the purpose of a plow? It's a tool that creates more: more seeds planted, more crops harvested, more food to go around. In places where nutrition is hard to come by, it's no exaggeration to say that a plow gives people more years of life. The plow—like many technologies, both ancient and modern—is about creating more of something and doing it more efficiently, so that more people can benefit.

Contrast that with lab-grown meat, one of the innovations I picked for this year's 10 Breakthrough Technologies list. Growing animal protein in a lab isn't about feeding more people. There's enough livestock to feed the world already, even as demand for meat goes up. Next-generation protein isn't about creating more—it's about making meat better. It lets us provide for a growing and wealthier world without contributing to deforestation or emitting methane. It also allows us to enjoy hamburgers without killing any animals.

Put another way, the plow improves our *quantity* of life, and lab-grown meat improves our *quality* of life. For most of human history, we've put most of our innovative capacity into the former. And our efforts have paid off: worldwide life expectancy rose from 34 years in 1913 to 60 in 1973 and has reached 71 today.

Read our interview with Bill Gates on page 56.

Because we're living longer, our focus is starting to shift toward well-being. This transformation is happening slowly. If you divide scientific breakthroughs into these two categories—things that improve quantity of life and things that improve quality of life—the 2009 list looks not so different from this year's. Like most forms of progress, the change is so gradual that it's hard to perceive. It's a matter of decades, not years—and I believe we're only at the midpoint of the transition.

To be clear, I don't think humanity will stop trying to extend life spans anytime soon. We're still far from a world where everyone everywhere lives to old age in perfect health, and it's going to take a lot of innovation to get us there. Plus, "quantity of life" and "quality of life" are not mutually exclusive. A malaria vaccine would both save lives and make life better for children who might otherwise have been left with developmental delays from the disease.

We've reached a point where we're tackling both ideas at once, and that's what makes this moment in history so interesting. If I had to predict what this list will look like a few years from now, I'd bet technologies that alleviate chronic disease will be a big theme. This won't just include new drugs (although I would love to see new treatments for diseases like Alzheimer's on the list). The innovations might look like a mechanical glove that helps a person with arthritis maintain flexibility, or an app that connects people experiencing major depressive episodes with the help they need.

If we could look even further out—let's say the list 20 years from now—I would hope to see technologies that center almost entirely on well-being. I think the brilliant minds of the future will focus on more metaphysical questions: How do we make people happier? How do we create meaningful connections? How do we help everyone live a fulfilling life?

I would love to see these questions shape the 2039 list, because it would mean that we've successfully fought back disease (and dealt with climate change). I can't imagine a greater sign of progress than that.

For now, though, the innovations driving change are a mix of things that extend life and things that make it better. My picks reflect both. Each one gives me a different reason to be optimistic for the future, and I hope they inspire you, too.

My selections include amazing new tools that will one day save lives, from simple blood tests that predict premature birth to toilets that destroy deadly pathogens. I'm equally excited by how other technologies on the list will improve our lives. Wearable health monitors like the wrist-based ECG will warn heart patients of impending problems, while others let diabetics not only track glucose levels but manage their disease. Advanced nuclear reactors could provide carbon-free, safe, secure energy to the world.

One of my choices even offers us a peek at a future where society's primary goal is personal fulfillment. Among many other applications, AI-driven personal agents might one day make your e-mail in-box more manageable—something that sounds trivial until you consider what possibilities open up when you have more free time.

The 30 minutes you used to spend reading e-mail could be spent doing other things. I know some people would use that time to get more work done—but I hope most would use it for pursuits like connecting with a friend over coffee, helping your child with homework, or even volunteering in your community.

That, I think, is a future worth working toward.

1/10 | By **WILL KNIGHT** | Illustration by Nico Ortega

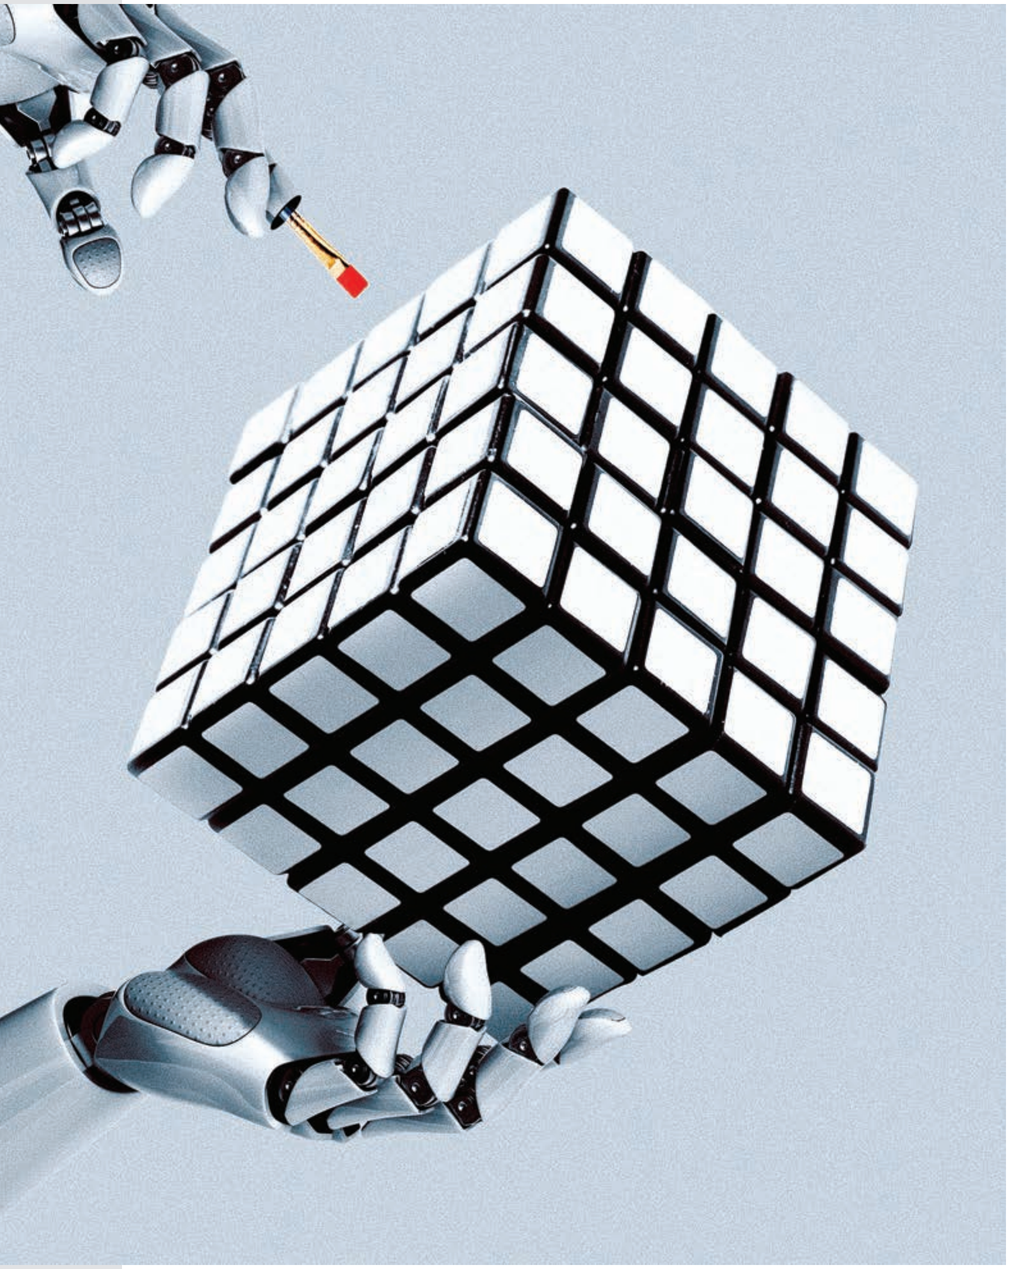

Robot dexterity

Robots are teaching themselves to handle the physical world.

Why it matters
-
If robots could learn to deal with the messiness of the real world, they could do many more tasks

Key players
-
OpenAI

Carnegie Mellon University

University of Michigan

UC Berkeley

Availability
-
3-5 years

For all the talk about machines taking jobs, industrial robots are still clumsy and inflexible. A robot can repeatedly pick up a component on an assembly line with amazing precision and without ever getting bored—but move the object half an inch, or replace it with something slightly different, and the machine will fumble ineptly or paw at thin air.

But while a robot can't yet be programmed to figure out how to grasp any object just by looking at it, as people do, it can now learn to manipulate the object on its own through virtual trial and error.

One such project is Dactyl, a robot that taught itself to flip a toy building block in its fingers. Dactyl, which comes from the San Francisco nonprofit OpenAI, consists of an off-the-shelf robot hand surrounded by an array of lights and cameras. Using what's known as reinforcement learning, neural-network software learns how to grasp and turn the block within a simulated environment before the hand tries it out for real. The software experiments, randomly at first, strengthening connections within the network over time as it gets closer to its goal.

If we can reliably employ this kind of learning, robots might eventually assemble our gadgets, load our dishwashers, and even help Grandma out of bed.

It usually isn't possible to transfer that type of virtual practice to the real world, because things like friction or the varied properties of different materials are so difficult to simulate. The OpenAI team got around this by adding randomness to the virtual training, giving the robot a proxy for the messiness of reality.

We'll need further breakthroughs for robots to master the advanced dexterity needed in a real warehouse or factory. But if researchers can reliably employ this kind of learning, robots might eventually assemble our gadgets, load our dishwashers, and even help Grandma out of bed.

2/10

By **LEIGH PHILLIPS**

Graphic by Tomi Um

Why it matters

-

Nuclear power is looking increasingly necessary in the effort to reduce carbon emissions and limit climate change

Key players

-

Terrestrial Energy

TerraPower

NuScale

General Fusion

Commonwealth Fusion Systems

Availability

-

New types of fission reactors could be widely available by the mid-2020s; fusion is more than a decade away

New-wave nuclear power

Advanced fusion and fission reactors are edging closer to reality.

New nuclear designs that have gained momentum in the past year are promising to make this power source safer and cheaper. Among them are generation IV fission reactors, an evolution of traditional designs; small modular reactors; and fusion reactors, a technology that has seemed eternally just out of reach. Developers of generation IV fission designs, such as Canada's Terrestrial Energy and Washington-based TerraPower, have entered into R&D partnerships with utilities, aiming for grid supply (somewhat optimistically, maybe) by the 2020s.

Small modular reactors typically produce in the tens of megawatts of power (for comparison, a traditional nuclear reactor produces around 1,000 MW). Companies like Oregon's NuScale say the miniaturized reactors can save money and reduce environmental and financial risks.

There has even been progress on fusion. Though no one expects delivery before 2030, companies like General Fusion and Commonwealth Fusion Systems, an MIT spinout, are making some headway. Many consider fusion a pipe dream, but because the reactors can't melt down and don't create long-lived, high-level waste, it should face much less public resistance than conventional nuclear. (Bill Gates is an investor in TerraPower and Commonwealth Fusion Systems.)

3/10

By **BONNIE ROCHMAN**

Predicting preemies

A simple blood test can predict if a pregnant woman is at risk of giving birth prematurely.

Our genetic material lives mostly inside our cells. But small amounts of “cell-free” DNA and RNA also float in our blood, often released by dying cells. In pregnant women, that cell-free material is an alphabet soup of nucleic acids from the fetus, the placenta, and the mother. Stephen Quake, a bioengineer at Stanford, has found a way to use that to tackle one of medicine’s most intractable problems: the roughly one in 10 babies born prematurely.

Free-floating DNA and RNA can yield information that previously required invasive ways of grabbing cells, such as taking a biopsy of a tumor or puncturing a pregnant woman’s belly to perform an amniocentesis. What’s changed is that it’s now easier to detect and sequence the small amounts of cell-free genetic material in the blood. In the last few years researchers have begun developing blood tests for cancer (by spotting the telltale DNA from tumor cells) and for prenatal screening of conditions like Down syndrome.

The tests for these conditions rely on looking for genetic mutations in the DNA. RNA, on the other hand, is the molecule that regulates gene expression—how much of a protein is produced from a gene. By sequencing the free-floating RNA in the mother’s blood, Quake can spot fluctuations in the expression of seven genes that he singles out as associated with preterm birth. That lets him identify women likely to deliver too early. Once alerted, doctors can take measures to stave off an early birth and give the child a better chance of survival.

The technology behind the blood test, Quake says, is quick, easy, and less than $10 a measurement. He and his collaborators have launched a startup, Akna Dx, to commercialize it.

Why it matters
-
15 million babies are born prematurely every year; it’s the leading cause of death for children under age five

Key players
-
Akna Dx

Availability
-
A test could be offered in doctor’s offices within five years

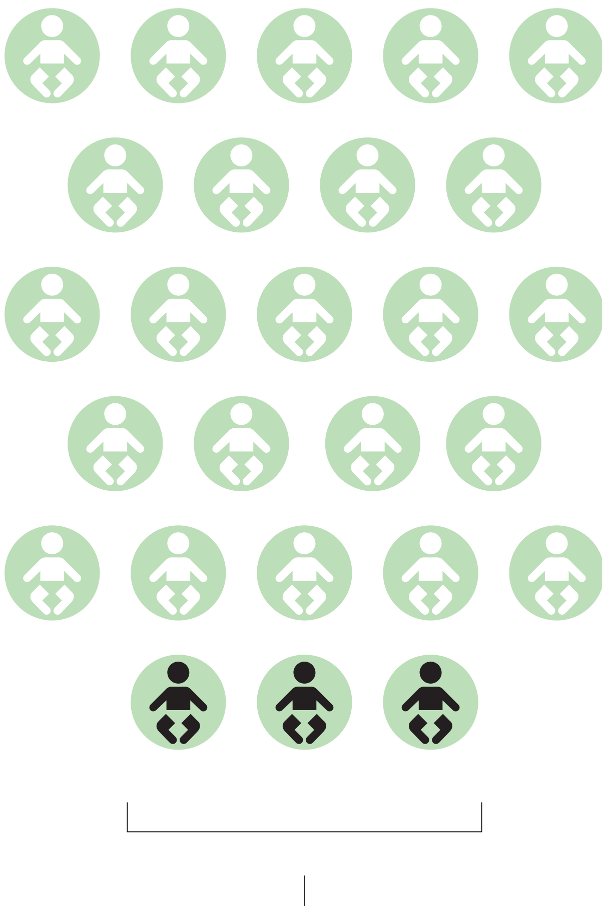

Every year, an estimated **15 million babies** are born preterm. (more than 1 in 10)

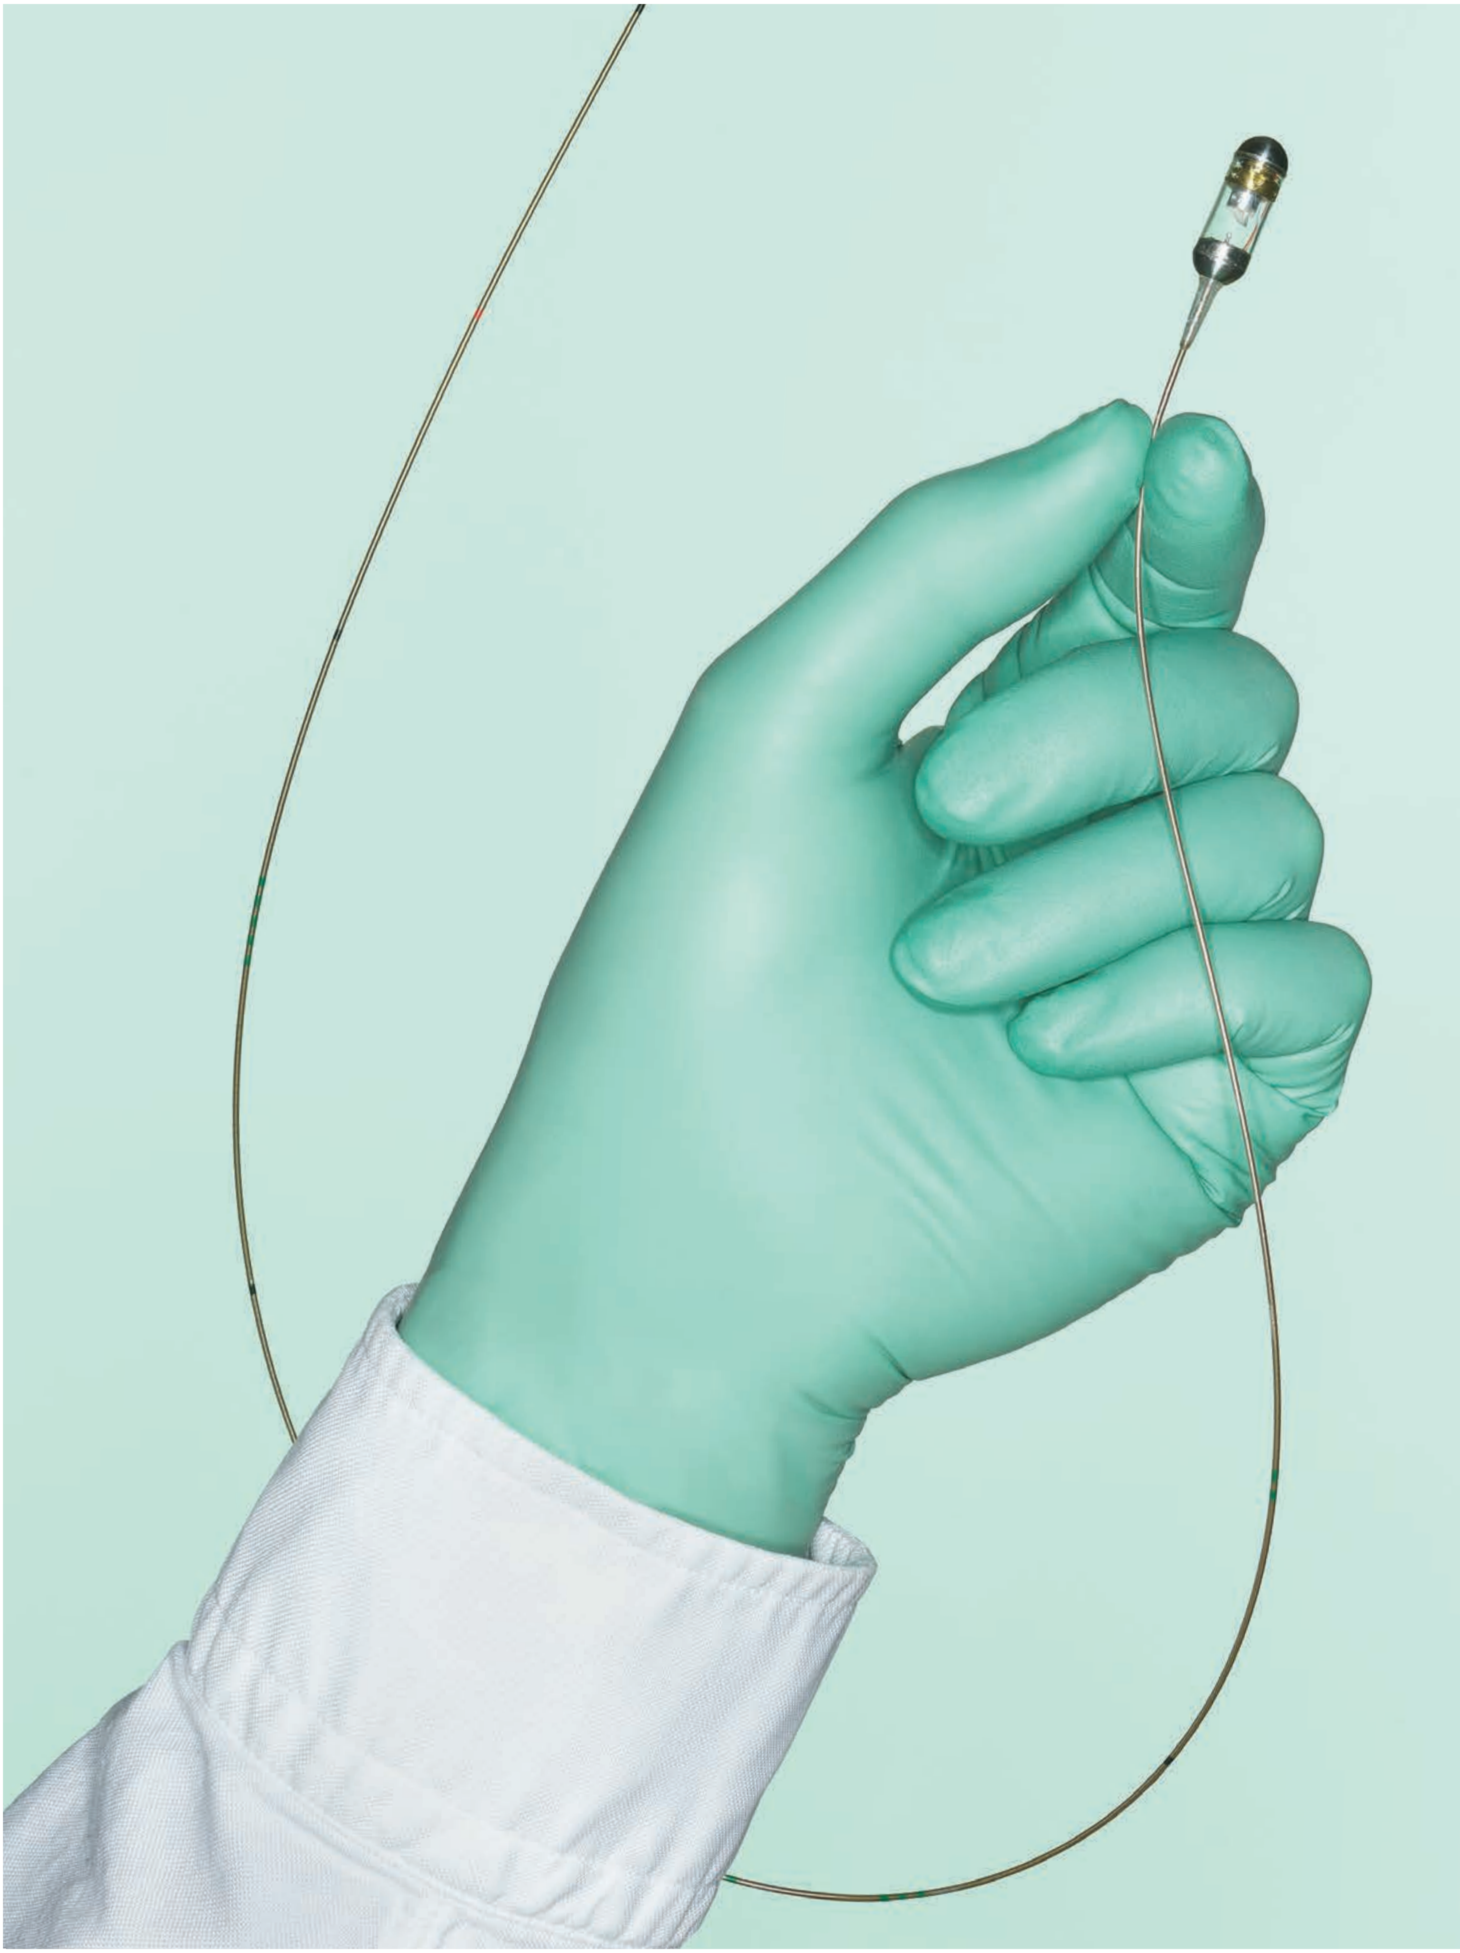

4/10

By COURTNEY HUMPHRIES

Photograph by Bruce Peterson

Gut probe in a pill

A small, swallowable device captures detailed images of the gut without anesthesia, even in infants and children.

Why it matters	Key players	Availability
The device makes it easier to screen for and study gut diseases, including one that keeps millions of children in poor countries from growing properly	Massachusetts General Hospital	Now used in adults; testing in infants begins in 2019

Environmental enteric dysfunction (EED) may be one of the costliest diseases you've never heard of. Marked by inflamed intestines that are leaky and absorb nutrients poorly, it's widespread in poor countries and is one reason why many people there are malnourished, have developmental delays, and never reach a normal height. No one knows exactly what causes EED and how it could be prevented or treated.

Practical screening to detect it would help medical workers know when to intervene and how. Therapies are already available for infants, but diagnosing and studying illnesses in the guts of such young children often requires anesthetizing them and inserting a tube called an endoscope down the throat. It's expensive, uncomfortable, and not practical in areas of the world where EED is prevalent.

So Guillermo Tearney, a pathologist and engineer at Massachusetts General Hospital (MGH) in Boston, is developing small devices that can be used to inspect the gut for signs of EED and even obtain tissue biopsies. Unlike endoscopes, they are simple to use at a primary care visit.

Tearney's swallowable capsules contain miniature microscopes. They're attached to a flexible string-like tether that provides power and light while sending images to a briefcase-like console with a monitor. This lets the health-care worker pause the capsule at points of interest and pull it out when finished, allowing it to be sterilized and reused. (Though it sounds gag-inducing, Tearney's team has developed a technique that they say doesn't cause discomfort.) It can also carry technologies that image the entire surface of the digestive tract at the resolution of a single cell or capture three-dimensional cross sections a couple of millimeters deep.

The technology has several applications; at MGH it's being used to screen for Barrett's esophagus, a precursor of esophageal cancer. For EED, Tearney's team has developed an even smaller version for use in infants who can't swallow a pill. It's been tested on adolescents in Pakistan, where EED is prevalent, and infant testing is planned for 2019.

The little probe will help researchers answer questions about EED's development—such as which cells it affects and whether bacteria are involved—and evaluate interventions and potential treatments.

By **BILL GATES**

Photograph by Bruce Peterson

What I'm reading

Whenever I want to understand something better, I pick up a book. Reading is my favorite way to learn about a new subject—whether it's global health, quantum computing, or world history. Here are 10 books that helped inform my choices for this year's list of 10 breakthrough technologies.

Enlightenment Now
by Steven Pinker

In my opening essay for this issue, I write about how innovation is increasingly aimed at improving quality of life. Pinker explains why in *Enlightenment Now* (which happens to be my favorite book). He looks at 15 different measures of progress to explain how and why the world is getting better.

Energy Myths and Realities
by Vaclav Smil

Smil convincingly argues that our present-day energy infrastructure will persist. He and I share a belief that nuclear power, which can use existing infrastructure while also reducing carbon emissions, will be an important electricity source for decades.

Life 3.0
by Max Tegmark

Anyone who wants to discuss how artificial intelligence is shaping the world should read this book. Tegmark, a physicist by training, takes a scientific approach. He doesn't spend a lot of time saying we should do this or that, and as a result, *Life 3.0* offers a terrific baseline of knowledge on the subject.

The Emperor of All Maladies
by Siddhartha Mukherjee

This Pulitzer Prize–winning "biography" of cancer is a beautifully told account of the progress made in fighting the disease over the last century. Some of the scientific advances that have resulted have led to other breakthroughs, like the vaccines included in this year's breakthrough technologies list.

Sustainable Energy—Without the Hot Air
by David MacKay

If you're interested in learning where energy comes from, how it is used, and what challenges are involved in switching to new sources, I can't recommend this book highly enough—and it will help you get more out of the next book on my list.

The Most Powerful Idea in the World
by William Rosen

For understanding how innovations change the world and evolve over time, Rosen's comprehensive history of the steam engine is as good a book as you will find.

Should We Eat Meat?
by Vaclav Smil

I'm a huge fan of everything Smil writes. He's skeptical that meat and dairy alternatives like those discussed in this issue will make a dent in global dietary habits. We might disagree on that particular point, but I think Smil has smart things to say about how to feed the world without destroying the planet.

Behind the Beautiful Forevers
by Katherine Boo

Boo's deeply reported narrative of life in a Mumbai slum might seem like an odd choice for a list of books about technology. But she offers perhaps the clearest look I've seen at the world's sanitation challenges. This one is essential reading for anyone hoping to reinvent the toilet.

I Contain Multitudes
by Ed Yong

I'm fascinated by microbes, and the human gut might hold the key to fixing all sorts of medical issues. I was particularly interested by Yong's account of how the bacteria that live in our digestive systems might be manipulated to prevent malnutrition.

Homo Deus
by Yuval Noah Harari

Harari describes a bleak future without sickness, hunger, and war—but where godlike elites and super-intelligent robots consider the rest of humanity to be superfluous. I'm more optimistic than he is about the chances of averting such a dystopia. If you're looking to tackle tomorrow's challenges, he offers some great food for thought.

By **ADAM PIORE**

Custom cancer vaccines

The treatment incites the body's natural defenses to destroy only cancer cells by identifying mutations unique to each tumor.

Scientists are on the cusp of commercializing the first personalized cancer vaccine. If it works as hoped, the vaccine, which triggers a person's immune system to identify a tumor by its unique mutations, could effectively shut down many types of cancers.

By using the body's natural defenses to selectively destroy only tumor cells, the vaccine, unlike conventional chemotherapies, limits damage to healthy cells. The attacking immune cells could also be vigilant in spotting any stray cancer cells after the initial treatment.

The possibility of such vaccines began to take shape in 2008, five years after the Human Genome Project was completed, when geneticists published the first sequence of a cancerous tumor cell.

Soon after, investigators began to compare the DNA of tumor cells with that of healthy cells—and other tumor cells. These studies confirmed that all cancer cells contain hundreds if not thousands of specific mutations, most of which are unique to each tumor.

A few years later, a German startup called BioNTech provided compelling evidence that a vaccine containing copies of these mutations could catalyze the body's immune system to produce T cells primed to seek out, attack, and destroy all cancer cells harboring them.

In December 2017, BioNTech began a large test of the vaccine in cancer patients, in collaboration with the biotech giant Genentech. The ongoing trial is targeting at least 10 solid cancers and aims to enroll upwards of 560 patients at sites around the globe.

The two companies are designing new manufacturing techniques to produce thousands of personally customized vaccines cheaply and quickly. That will be tricky because creating the vaccine involves performing a biopsy on the patient's tumor, sequencing and analyzing its DNA, and rushing that information to the production site. Once produced, the vaccine needs to be promptly delivered to the hospital; delays could be deadly.

Why it matters
-
Conventional chemotherapies take a heavy toll on healthy cells and aren't always effective against tumors

Key players
-
BioNTech
Genentech

Availability
-
In human testing

To create the vaccine:

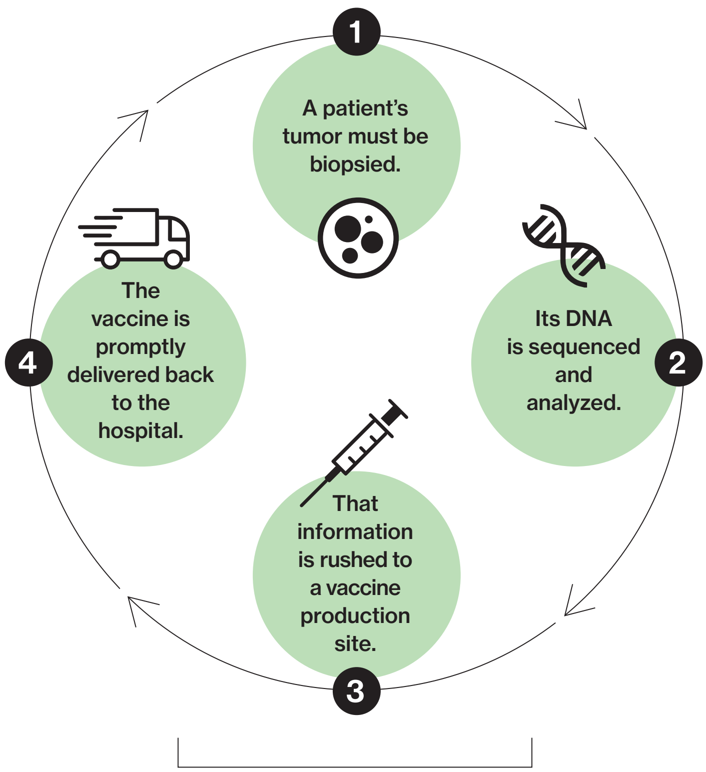

Any delay could be deadly.

By **THE EDITORS**

10 grand challenges

These are big problems that new technologies might solve or fundamental questions they might answer. Some might be solved one day, while others may remain unconquerable. None are easy, but all of them, we think, are incredibly important.

Universal flu vaccine

Pandemic flu is rare but deadly. At least 50 million people died in the 1918 pandemic of H1N1 flu. More recently, about a million people died in the 1957-'58 and 1968 pandemics, while something like half a million died in a 2009 recurrence of H1N1. The recent death tolls are lower in part because the viruses were milder strains. We might not be so lucky next time—a particularly potent strain of the virus could replicate too quickly for any tailor-made vaccine to effectively fight it. A universal flu vaccine that protected not only against the relatively less harmful variants but also against a catastrophic once-in-a-century outbreak is a crucial public health challenge.

Carbon sequestration

Cutting greenhouse-gas emissions alone won't be enough to prevent sharp increases in global temperatures. We'll also need to remove vast amounts of carbon dioxide from the atmosphere, which not only would be incredibly expensive but would present us with the thorny problem of what to do with all that CO_2 (see "Is carbon removal crazy or critical? Yes," on page 28). A growing number of startups are exploring ways of recycling carbon dioxide into products, including synthetic fuels, polymers, carbon fiber, and concrete. That's promising, but what we'll really need is a cheap way to permanently store the billions of tons of carbon dioxide that we might have to pull out of the atmosphere.

Dementia treatment

More than one in 10 Americans over the age of 65 has Alzheimer's; a third of those over 85 do. As people's lifespans lengthen, the number of people living with the disease—in the US and around the world—is likely to skyrocket. Alzheimer's remains poorly understood: conclusive diagnoses are possible only after death, and even then, doctors debate the distinction between Alzheimer's and other forms of dementia. However, advances in neuroscience and genetics are beginning to shed more light. That understanding is providing clues to how it might be possible to slow or even shut down the devastating effects of the condition.

Grid-scale energy storage

Renewable energy sources like wind and solar are becoming cheap and more widely deployed, but they don't generate electricity when the sun's not shining or wind isn't blowing. That limits how much power these sources can supply, and how quickly we can move away from steady sources like coal and natural gas. The cost of building enough batteries to back up entire grids for the days when renewable generation flags would be astronomical. Various scientists and startups are working to develop cheaper forms of grid-scale storage that can last for longer periods, including flow batteries or tanks of molten salt. Either way, we desperately need a cheaper and more efficient way to store vast amounts of electricity.

Ocean clean-up

Billions of tiny pieces of plastic—so-called "microplastics"—are now floating throughout the world's oceans. Much of this waste comes from bags or straws that have been broken up over time. It's poisoning birds, fish, and humans. Researchers fear that the effects on both human health and the environment will be profound, and it may take centuries to clean up the hundreds of millions of tons of plastic that have accumulated over the decades. Because the pollution is so diffuse, it's difficult to clean up, and while there are prototype methods for tackling the massive oceanic garbage patches, there is no solution for coasts, seas, and waterways.

Graphics by Tomi Um

Energy-efficient desalination

There is about 50 times as much salt water on earth as there is fresh water. As the world's population grows and climate change intensifies droughts, the need for fresh water is going to grow more acute. Israel has built the world's biggest reverse-osmosis desalination facilities and now gets most of its household water from the sea, but that method is too energy intensive to be practical worldwide. New types of membranes might help; electrochemical techniques may also help to make brackish water useful for irrigation. As far as climate-change adaptation technologies go, creating drinking water from the ocean ought to be a top priority.

Earthquake prediction

Over 100,000 people died in the 2010 Haiti earthquake, and the 2004 Indian Ocean tsunami—triggered by one of the most powerful earthquakes ever recorded—killed nearly a quarter of a million people in Indonesia, Sri Lanka, India, and elsewhere. We can predict hurricanes days and sometimes weeks in advance, but earthquakes still come as a surprise. Predicting earthquakes with some confidence over the medium term would allow planners to figure out durable solutions. At least giving a few hours' warning would allow people to evacuate unsafe areas, and could save millions of lives.

Safe driverless car

Autonomous vehicles have been tested for millions of miles on public roads. Pilot programs for delivery and taxi services are under way in places like the suburbs of Phoenix. But driverless cars still aren't ready to take over roads in general. They have trouble handling chaotic traffic, and difficulty with weather conditions like snow and fog. If they can be made reliably safe, they might allow a wholesale reimagining of transportation. Traffic jams might be eliminated, and cities could be transformed as parking lots give way to new developments. Above all, self-driving cars, if widely deployed, are expected to eliminate most of the 1.25 million deaths a year caused by traffic accidents.

We can predict hurricanes days and sometimes weeks in advance, but earthquakes still come as a surprise. Predicting them with confidence could save millions of lives.

Embodied AI

Last fall a video of Atlas, designed by Boston Dynamics, swept the internet. It showed the robot jumping up steps like a commando. This came only two years after AlphaGo beat the world's best Go player. Atlas can't play Go (it is embodied, but not intelligent), and AlphaGo can't run (it's intelligent, in its own way, but lacks a body). So what happens if you put AlphaGo's mind in Atlas's body? Many researchers say true general artificial intelligence might depend on an ability to relate internal computational processes to real things in the physical world, and that an AI would acquire that ability by learning to interact with the physical world as people and animals do.

Brain decoding

Our brains remain a deep mystery to neuroscientists. Everything we think and remember, and all our movements, must somehow be coded in the billions of neurons in our heads. But what is that code? There are still many unknowns and puzzles in understanding the way our brains store and communicate our thoughts. Cracking that code could lead to breakthroughs in how we treat mental disorders like schizophrenia and autism. It might allow us to improve direct interfaces that communicate directly from our brains to computers, or even to other people—a life-changing development for people who are paralyzed by injury or degenerative disease.

6/10

By **MARKKUS ROVITO**

Photograph by Bruce Peterson

STYLING BY MONICA MARIANO/ENNIS

The cow-free burger

Both lab-grown and plant-based alternatives approximate the taste and nutritional value of real meat without the environmental devastation.

Why it matters
-
Livestock production causes catastrophic deforestation, water pollution, and greenhouse-gas emissions

Key players
-
Beyond Meat

Impossible Foods

Availability
-
Plant-based now; lab-grown around 2020

The UN expects the world to have 9.8 billion people by 2050. And those people are getting richer. Neither trend bodes well for climate change—especially because as people escape poverty, they tend to eat more meat.

By that date, according to the predictions, humans will consume 70% more meat than they did in 2005. And it turns out that raising animals for human consumption is among the worst things we do to the environment.

Depending on the animal, producing a pound of meat protein with Western industrialized methods requires 4 to 25 times more water, 6 to 17 times more land, and 6 to 20 times more fossil fuels than producing a pound of plant protein.

The problem is that people aren't likely to stop eating meat anytime soon. Which means lab-grown and plant-based alternatives might be the best way to limit the destruction.

Making lab-grown meat involves extracting muscle tissue from animals and growing it in bioreactors. The end product looks much like what you'd get from an animal, although researchers are still working on the taste. Researchers at Maastricht University in the Netherlands, who are working to produce lab-grown meat at scale, believe that by next year a lab-grown burger could cost no more than a hamburger made from a cow. One drawback of lab-grown meat is that the environmental benefits are still sketchy at best—a recent World Economic Forum report says the emissions from lab-grown meat would be only around 7% less than emissions from beef production.

The better environmental case can be made for plant-based meats from companies like Beyond Meat and Impossible Foods (Bill Gates is an investor in both companies), which use pea proteins, soy, wheat, potatoes, and plant oils to mimic the texture and taste of animal meat. Beyond Meat has a new 26,000-square-foot (2,400-square-meter) plant in California and has already sold upwards of 25 million burgers from 30,000 stores and restaurants. According to an analysis by the Center for Sustainable Systems at the University of Michigan, a Beyond Meat patty would probably generate 90% less greenhouse-gas emissions than a conventional burger made from a cow.

By THE EDITORS

The low-tech 10

Technologies don't have to be cutting edge to make a profound difference in people's lives.

Oral rehydration salts

By the early 1990s, diarrheal diseases were killing some 5 million children under the age of five every year. That number is down to about 1.5 million, thanks to oral rehydration salts—a mixture of salt and sugar that can be dissolved in water and administered at home. Zinc is sometimes added to the mix to reduce the severity and duration of diarrhea. This simple innovation has perhaps saved more lives at lower cost than any other.

Cheap, low-power irrigation

Irrigation accounts for the bulk of freshwater use in most countries—something like three quarters of the total. Drip irrigation uses half as much water as conventional irrigation and is half again as productive. But it's expensive and usually requires electrical power. The GEAR lab at MIT has developed low-pressure solar-powered drip irrigation systems that can deliver the benefits at much lower cost.

DC-power microgrid

Solar cells can provide cheap, decentralized electricity. But if you're plugging them into conventional devices on a normal household grid, there's a lot of overhead involved in converting the direct current they produce into alternating current and back again. A well-designed small DC network can save a substantial amount of energy by eliminating this need.

Better woodstoves

Deforestation is a major problem in much of the developing world, as is the harm to human health that comes from breathing in the particulate matter in smoke from woodstoves. Better-designed stoves like the Berkeley-Darfur stove use only half as much fuel to cook a comparable amount of food, and they cut the particulate emissions in half as well.

Simple, effective water filters

Hundreds of millions of people around the world lack access to safe water. Simple, cheap water filters use ash combined with silver nanoparticles to filter out impurities and pathogens; they have improved the lives of hundreds of thousands.

Hippo roller

Hundreds of millions of people, usually women, have to walk every day to get enough water for their basic needs and transport it home in buckets. The Hippo roller is a heavy-duty plastic barrel that can be flipped on its side and rolled home, via an attached handle, over rough terrain.

Jet injections

Vaccines are crucial for public health. But in the developing world, distributing the vaccine to where it's needed is only part of the problem. How do you administer it in a place where sterile needles might be scarce? One fix is a jet injector, a decades-old invention that can send a high-pressure, directed stream of fluid through the skin.

Paper microscopes

Microscopes are crucial for diagnosing infectious disease. But in some ways they're the worst possible device—heavy, expensive, and hard to maintain. Paper microscopes, also known as foldscopes, contain all the crucial parts within one foldable sheet of paper. They can be optimized for different diseases and cost less than a dollar.

Disaster communications system

Cell phones are common even in poor countries, but when a natural disaster strikes, the communications networks these devices rely upon can fail. Developed in Chile, SiE is a system that encodes text into high-frequency audio tones that can be distributed over broadcast radio waves and received on any smartphone without requiring any internet infrastructure. An app on the phone listens for these tones and transforms them into a text message.

Portable malaria screener

Malaria kills 3,000 children a day. Quick diagnosis and treatment is crucial, but that typically requires a microscope and a reliable technician to analyze blood samples. A quicker, simpler system developed last year at the University of Southern California is portable and detects levels of hemozoin, a by-product created by the malaria parasite, which reveals how far the disease has progressed.

By **JAMES TEMPLE**

Illustration by Nico Ortega

Carbon dioxide catcher

Practical and affordable ways to capture carbon dioxide from the air can soak up excess greenhouse-gas emissions.

Why it matters
-
Removing CO_2 from the atmosphere might be one of the last viable ways to stop catastrophic climate change

Key players
-
Carbon Engineering

Climeworks

Global Thermostat

Availability
-
5-10 years

Even if we slow carbon dioxide emissions, the warming effect of the greenhouse gas can persist for thousands of years. To prevent a dangerous rise in temperatures, the UN's climate panel now concludes, the world will need to remove as much as 1 trillion tons of carbon dioxide from the atmosphere this century.

In a surprise finding last summer, Harvard climate scientist David Keith calculated that machines could, in theory, pull this off for less than $100 a ton, through an approach known as direct air capture. That's an order of magnitude cheaper than earlier estimates that led many scientists to dismiss the technology as far too expensive—though it will still take years for costs to fall to anywhere near that level.

But once you capture the carbon, you still need to figure out what to do with it.

Carbon Engineering, the Canadian startup Keith cofounded in 2009, plans to expand its pilot plant to ramp up production of its synthetic fuels, using the captured carbon dioxide as a key ingredient. (Bill Gates is an investor in Carbon Engineering.)

Pulling CO_2 from the air is a difficult way of dealing with climate change, but we're running out of options.

Zurich-based Climeworks's direct air capture plant in Italy will produce methane from captured carbon dioxide and hydrogen, while a second plant in Switzerland will sell carbon dioxide to the soft-drinks industry. So will Global Thermostat of New York, which finished constructing its first commercial plant in Alabama last year.

Still, if it's used in synthetic fuels or sodas, the carbon dioxide will mostly end up back in the atmosphere. The ultimate goal is to lock greenhouse gases away forever. Some could be nested within products like carbon fiber, polymers, or concrete, but far more will simply need to be buried underground, a costly job that no business model seems likely to support.

In fact, pulling CO_2 out of the air is, from an engineering perspective, one of the most difficult and expensive ways of dealing with climate change. But given how slowly we're reducing emissions, there are no good options left.

8/10

By **KAREN HAO**

Graphic by Tomi Um

Why it matters
-
Wearable ECGs can warn of potentially life-threatening cardiac problems such as atrial fibrillation

Key players
-
Apple
AliveCor
Withings

Availability
-
Now

An ECG on your wrist

Regulatory approval and technological advances are making it easier for people to continuously monitor their hearts with wearable devices.

Fitness trackers aren't serious medical devices. An intense workout or loose band can mess with the sensors that read your pulse. But an electrocardiogram—the kind doctors use to diagnose abnormalities before they cause a stroke or heart attack—requires a visit to a clinic, and people often fail to take the test in time.

ECG-enabled smart watches, made possible by new regulations and innovations in hardware and software, offer the convenience of a wearable device with something closer to the precision of a medical one.

An Apple Watch–compatible band from Silicon Valley startup AliveCor that can detect atrial fibrillation, a frequent cause of blood clots and stroke, received clearance from the FDA in 2017. Last year, Apple released its own FDA-cleared ECG feature, embedded in the watch itself. The health-device company Withings also announced plans for an ECG-equipped watch shortly after.

Current wearables still employ only a single sensor, whereas a real ECG has 12. And no wearable can yet detect a heart attack as it's happening.

But this might change soon. Last fall, AliveCor presented preliminary results to the American Heart Association on an app and two-sensor system that can detect a certain type of heart attack.

9/10

By **ERIN WINICK**

Sanitation without sewers

Energy-efficient toilets can operate without a sewer system and treat waste on the spot.

About 2.3 billion people don't have good sanitation. The lack of proper toilets encourages people to dump fecal matter into nearby ponds and streams, spreading bacteria, viruses, and parasites that can cause diarrhea and cholera. Diarrhea causes one in nine child deaths worldwide.

Now researchers are working to build a new kind of toilet that's cheap enough for the developing world and can not only dispose of waste but treat it as well.

In 2011 Bill Gates created what was essentially the X Prize in this area—the Reinvent the Toilet Challenge. Since the contest's launch, several teams have put prototypes in the field. All process the waste locally, so there's no need for large amounts of water to carry it to a distant treatment plant.

Most of the prototypes are self-contained and don't need sewers, but they look like traditional toilets housed in small buildings or storage containers. The NEWgenerator toilet, designed at the University of South Florida, filters out pollutants with an anaerobic membrane, which has pores smaller than bacteria and viruses. Another project, from Connecticut-based Biomass Controls, is a refinery the size of a shipping container; it heats the waste to produce a carbon-rich material that can, among other things, fertilize soil.

One drawback is that the toilets don't work at every scale. The Biomass Controls product, for example, is designed primarily for tens of thousands of users per day, which makes it less well suited for smaller villages. Another system, developed at Duke University, is meant to be used only by a few nearby homes.

So the challenge now is to make these toilets cheaper and more adaptable to communities of different sizes. "It's great to build one or two units," says Daniel Yeh, an associate professor at the University of South Florida, who led the NEWgenerator team. "But to really have the technology impact the world, the only way to do that is mass-produce the units."

Why it matters
-
2.3 billion people lack safe sanitation, and many die as a result

Key players
-
Duke University

University of South Florida

Biomass Controls

California Institute of Technology

Availability
-
1-2 years

The number of people who

Still do not have basic sanitation facilities such as toilets or latrines:

2.3 BILLION

Are thought to consume food irrigated by wastewater:

.75 BILLION

Die in low- and middle-income countries each year as a result of inadequate water, sanitation, and hygiene:

842,000

Still defecate in the open, for example in street gutters, behind bushes, or into open bodies of water:

892 MILLION

SOURCE: WHO

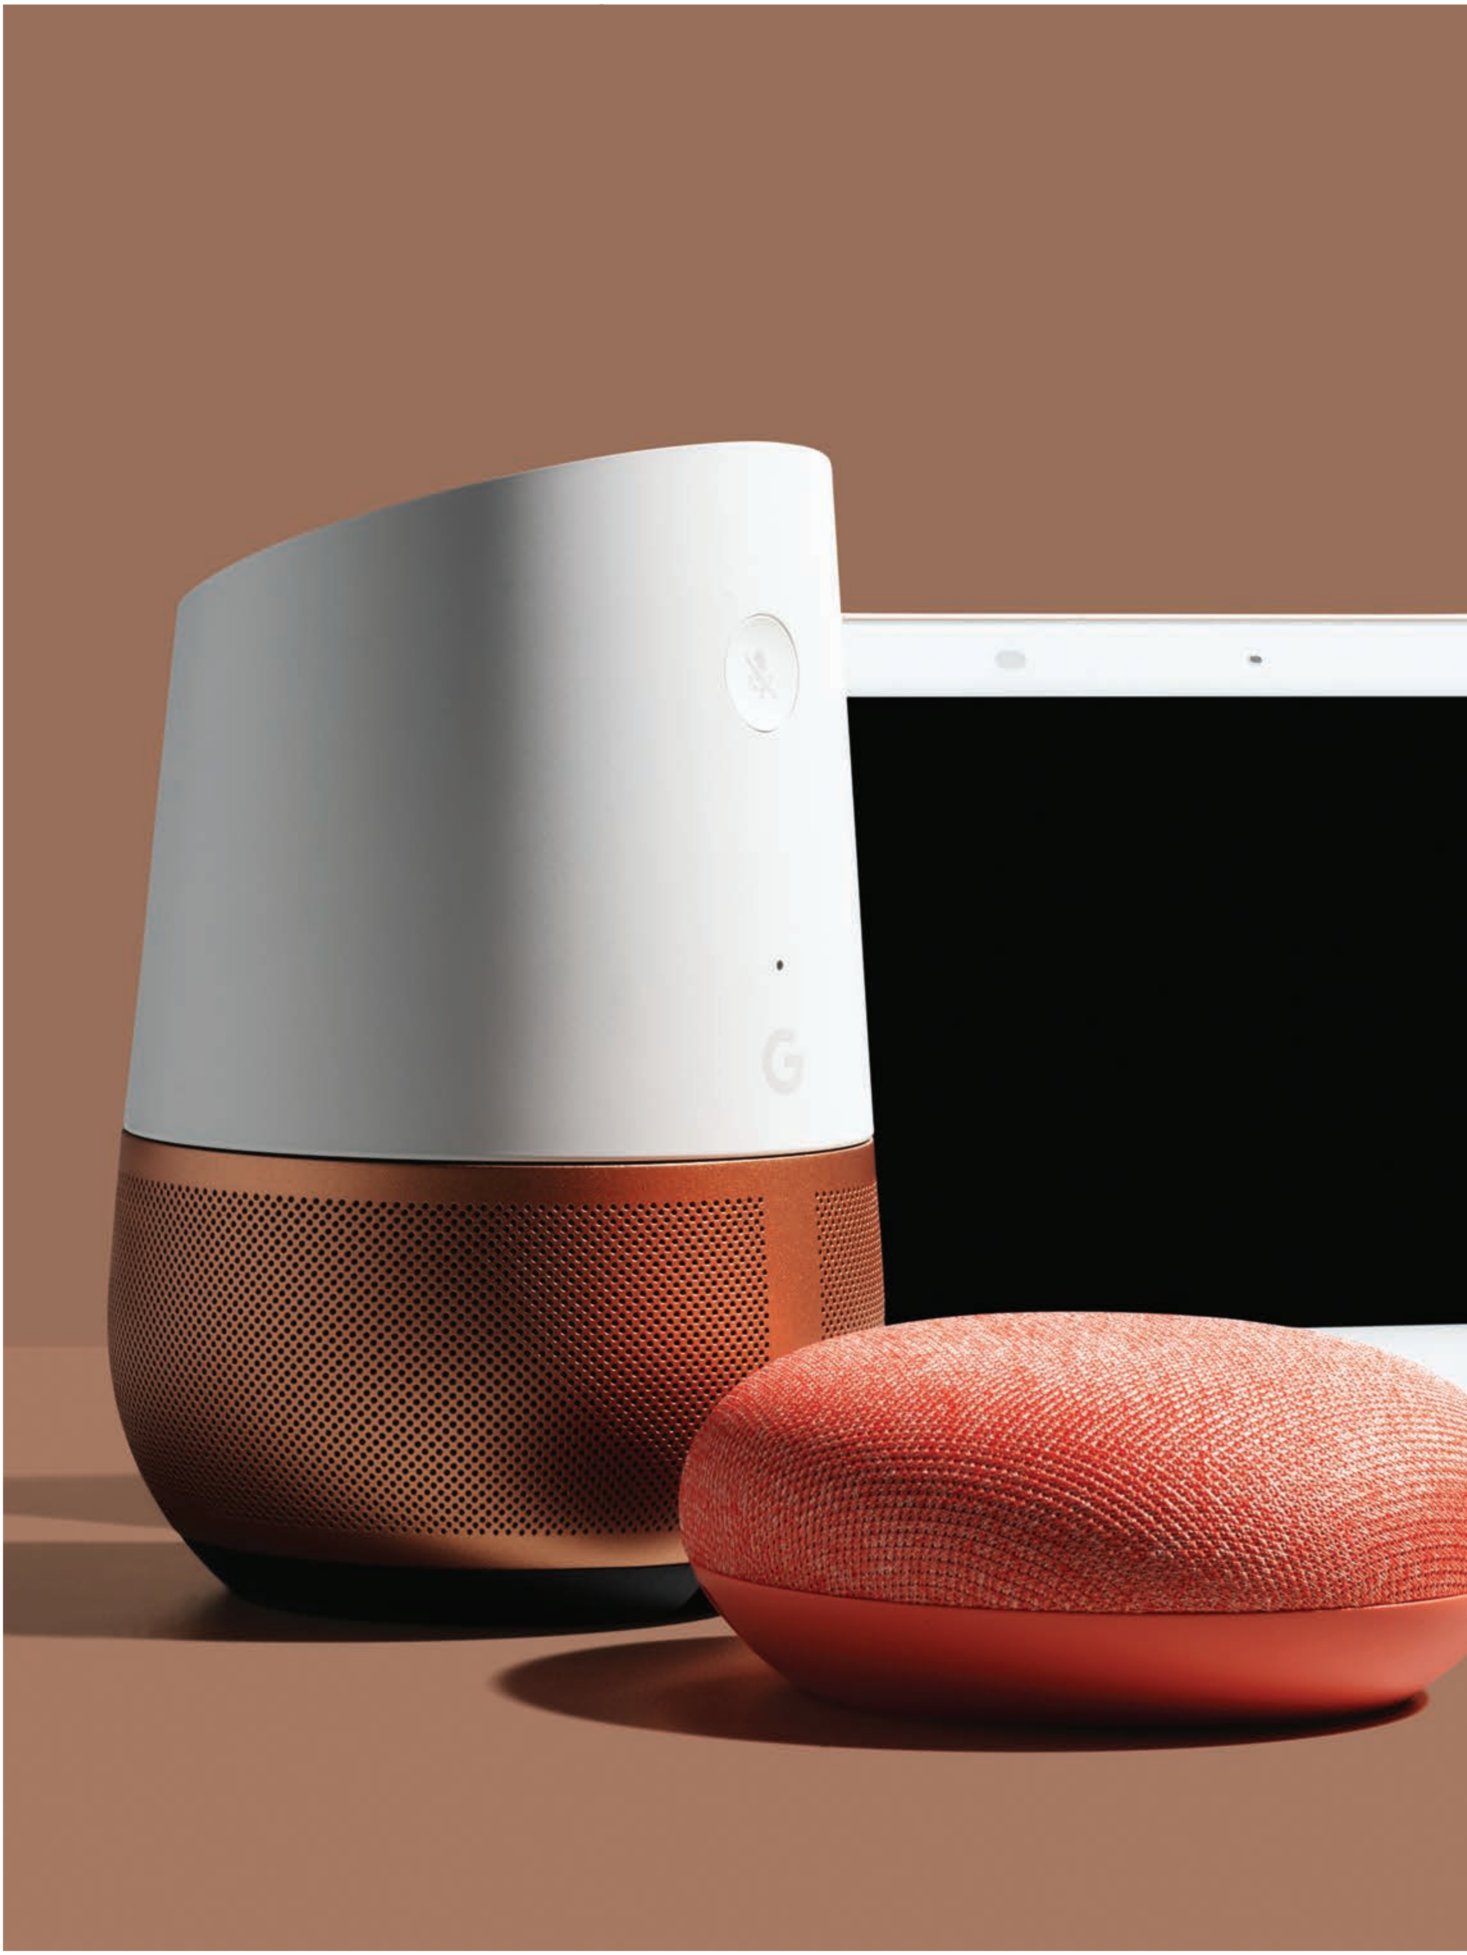
G

10/10 | By **KAREN HAO** | Photograph by Bruce Peterson

Smooth-talking AI assistants

Why it matters
-
AI assistants can now perform conversation-based tasks like booking a restaurant reservation or coordinating a package drop-off rather than just obey simple commands

Key players
-
Google

Alibaba

Amazon

Availability
-
1-2 years

New techniques that capture semantic relationships between words are making machines better at understanding natural language.

We're used to AI assistants—Alexa playing music in the living room, Siri setting alarms on your phone—but they haven't really lived up to their alleged smarts. They were supposed to have simplified our lives, but they've barely made a dent. They recognize only a narrow range of directives and are easily tripped up by deviations.

But some recent advances are about to expand your digital assistant's repertoire. In June 2018, researchers at OpenAI developed a technique that trains an AI on unlabeled text to avoid the expense and time of categorizing and tagging all the data manually. A few months later, a team at Google unveiled a system called BERT that learned how to predict missing words by studying millions of sentences. In a multiple-choice test, it did as well as humans at filling in gaps.

These improvements, coupled with better speech synthesis, are letting us move from giving AI assistants simple commands to having conversations with them. They'll be able to deal with daily minutiae like taking meeting notes, finding information, or shopping online.

Some are already here. Google Duplex, the eerily human-like upgrade of Google Assistant, can pick up your calls to screen for spammers and telemarketers. It can also make calls for you to schedule restaurant reservations or salon appointments.

In China, consumers are getting used to Alibaba's AliMe, which coordinates package deliveries over the phone and haggles about the price of goods over chat.

But while AI programs have gotten better at figuring out what you want, they still can't understand a sentence. Lines are scripted or generated statistically, reflecting how hard it is to imbue machines with true language understanding. Once we cross that hurdle, we'll see yet another evolution, perhaps from logistics coordinator to babysitter, teacher—or even friend?

IS CARBON REMOVAL

CRAZY OR CRITICAL?

Photographs by Spencer Lowell

By James Temple

YES.

The big metal container in Klaus Lackner's lab doesn't look as if it could save the planet. It most closely resembles a dumpster— which it sort of is.

As Lackner looks on, hands in the pockets of his pressed khakis, the machine begins to transform. Three mattress-shaped metal frames rise from the guts of the receptacle, unfolding like an accordion as they stretch toward the ceiling.

Each frame contains hundreds of white polymer strips filled with resins that bind with carbon dioxide molecules. The strips form a kind of sail, designed to snatch the greenhouse gas out of the air as wind blows through the contraption.

Crucially, that same material releases the carbon dioxide when wet. To make that happen, Lackner's device retracts its frames into their container, which then fills with water. The gas can next be collected and put to other uses, and the process can begin again.

Lackner and his colleagues at Arizona State University's Center for Negative Carbon Emissions have built a simple machine with a grand purpose: capturing and recycling carbon dioxide to ease the effects of climate change. He envisions forests of them stretching across the countryside, sucking up billions of tons of it from the atmosphere.

Lackner, 66, with receding silver hair, has now been working on the problem for two decades. In 1999, as a particle physicist at Los Alamos National Laboratory, he wrote the first scientific paper exploring the feasibility of combating climate change by pulling carbon dioxide out of the air. His was a lonely voice for years. But a growing crowd has come around to his thinking as the world struggles to slash climate emissions fast enough to prevent catastrophic warming. Lackner's work has helped inspire a handful of direct-air-capture startups, including one of his own, and a growing body of scientific literature. "It's hard to think of another field that is so much the product of a single person's thinking and advocacy," says David Keith, a Harvard professor who cofounded another of those startups, Carbon Engineering. "Klaus was pivotal in making the argument that [direct air capture] could be developed at a scale relevant to the carbon-climate problem."

The carbon-trapping materials work in various forms, including a grass-like structure used to fertilize greenhouses (previous pages).

The latest prototype (right) unfolds to grab carbon from the air. Klaus Lackner (next page) pioneered the field of direct air capture.

No one, including Lackner, really knows whether the scheme will work. The chemistry is easy enough. But can we really construct anywhere near enough carbon removal machines to make a dent in climate change? Who will pay for them? And what are we going to do with all the carbon dioxide they collect?

Lackner readily acknowledges the unknowns but believes that the cheaper the process gets, the more feasible it becomes. "If I tell you, 'You could solve the carbon problem for $1,000 a ton,' we will say, 'Climate change is a hoax,'" Lackner says. "But if it's $5 a ton, or $1 a ton, we'll say, 'Why haven't we fixed it yet?'"

Narrowing our options

The concentration of carbon dioxide in the atmosphere is approaching 410

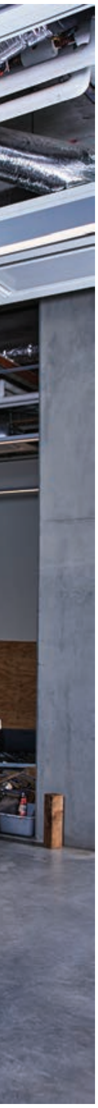

parts per million. That has already driven global temperatures nearly 1 °C above preindustrial levels and intensified droughts, wildfires, and other natural disasters. Those dangers will only compound as emissions continue to rise.

The latest assessment from the UN's Intergovernmental Panel on Climate Change found that there's no way to limit or return global warming to 1.5 °C without removing somewhere between 100 billion and a trillion metric tons of carbon dioxide by the end of the century. On the high end, that means reversing nearly three decades of global emissions at the current rate.

There are a handful of ways to draw carbon dioxide out of the atmosphere. They include planting lots of trees, restoring grasslands and other areas that naturally hold carbon in soils, and using carbon dioxide–sucking plants and other forms of biomass as a fuel source but capturing any emissions when they're used (a process known as bio-energy with carbon capture and storage).

But a report from the US National Academies in October found that these approaches alone probably won't be enough to prevent 2 °C of warming—at least, not if we want to eat. That's because the amount of land required to capture that much carbon dioxide would come at the cost of a huge amount of agricultural food production.

The appeal of direct-air-capture devices like the ones Lackner and others are developing is that they can suck out the same amount of carbon dioxide on far less land. The big problem is that right now it's much cheaper to plant a tree. At the current cost of around $600 per ton, capturing a trillion tons would run some $600 trillion, more than seven times the world's annual GDP.

In a paper last summer, Harvard's Keith calculated that the direct-air-capture system he helped design could eventually cost less than $100 a ton at full scale. Carbon Engineering, based in British Columbia, is in the process of expanding its pilot plant to increase production of synthetic fuels, created by combining the captured carbon dioxide with hydrogen. These, in turn, will be converted into forms of diesel and jet fuel that are considered carbon neutral, since they don't require digging up additional fossil fuels.

If Keith's method can capture carbon dioxide for $100 a ton, these synthetic fuels could be sold profitably in markets with public policy support, such as California, with its renewable-fuel standards, or the European Union, under its updated Renewable Energy Directive. The hope is that these kinds of early opportunities will help scale up the technology, drive down costs further, and open additional markets.

Other startups, including Switzerland-based Climeworks and Global Thermostat of New York, think they can achieve similar or even lower costs. They are exploring markets like the soda industry and greenhouses, which use air enriched with carbon dioxide to fertilize plants.

However, selling carbon dioxide isn't an easy proposition.

Global demand is relatively small: on the order of a few hundred million tons per year, a fraction of the tens of billions that eventually need to be removed annually, according to the National Academies study. Moreover, most of that demand is for enhanced oil recovery, a technique that forces compressed carbon dioxide into wells to free up the last drips of oil, which only makes the climate problem worse.

A critical question for the carbon-capture startups is how much the market for carbon dioxide could grow. Dozens of businesses are exploring new ways of putting it to work. They include

"SO THE IDEA THAT WE'RE GOING TO GET TO NEGATIVE CIVILIZATION-SCALE EMISSIONS THROUGH AIR CAPTURE, TO ME, JUST SEEMS LIKE A FANTASY."

California-based Opus12, which is using carbon dioxide to produce chemicals and polymers, and CarbonCure of Nova Scotia, which is working with more than 100 concrete manufacturers to convert carbon dioxide into calcium carbonate that gets trapped in the concrete as it sets.

A 2016 report by the Global CO_2 Initiative estimated that the market for products that could use carbon dioxide—including liquid fuels, polymers, methanol, and concrete—could reach $800 billion by 2030. Those industries could put to use some 7 billion metric tons per year—about 15% of annual global emissions.

Such projections are extremely optimistic, though. And even if such a vast transformation of multiple sectors actually occurs, it will still leave huge amounts of captured carbon dioxide that will need to be permanently stored underground.

That's only going to happen if society decides to pay for it, and some are skeptical we ever will. Capturing carbon dioxide out of the air—which means plucking a single molecule from amid nearly 2,500 others—is one of the most energy-intensive and expensive ways we could dream up of grappling with climate change. "Direct air capture is more expensive than avoiding emissions, but right now we're not even willing to spend the additional money to do that," says Ken Caldeira, a climate scientist at the Carnegie Institution. "So the idea that we're going to get to negative civilization-scale emissions through air capture, to me, just seems like a fantasy."

Lackner peers through an early model of an air-capture device, with the carbon-trapping materials shaped into a grid.

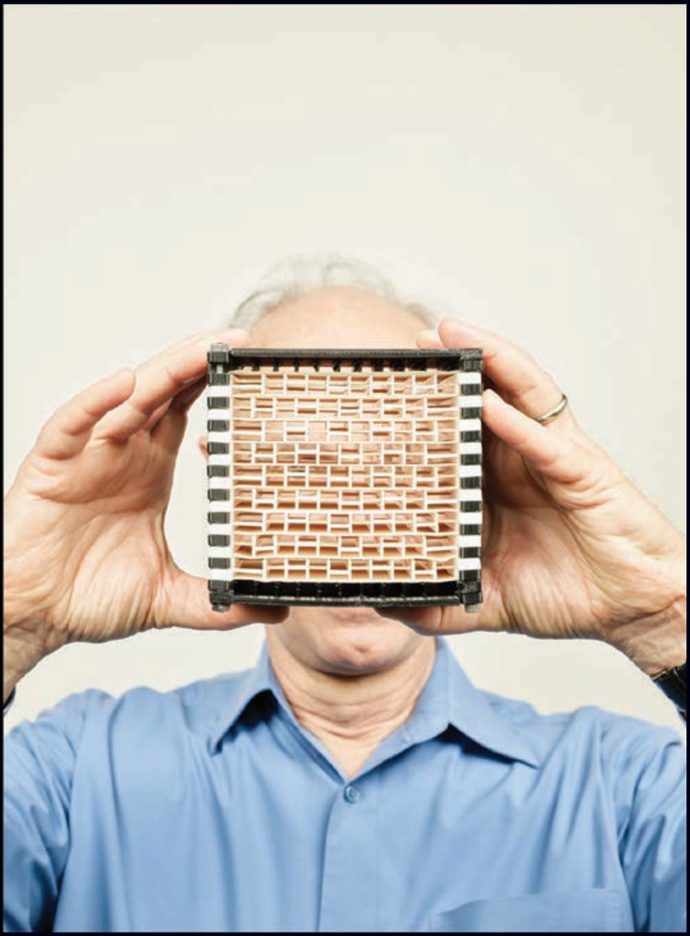

Robot-making robots

On a summer night in 1992, while Lackner was a researcher at Los Alamos National Laboratory, he and a fellow particle physicist were having a beer and complaining about the lack of big, bold ideas in science. One or two drinks later, they had one of their own: What would become possible if machines could build machines? How big and fast could you manufacture things?

They quickly realized that the only way the scheme would work is if you designed robots that dug up all their own raw materials from dirt, constructed solar panels to power the process—and made ever more copies of themselves.

The next morning, Lackner and his friend, Christopher Wendt of the University of Wisconsin–Madison, decided they had an idea worth exploring. They eventually published a paper working out the math and exploring several applications, including self-replicating robots that could

GUTTER CREDIT HERE

capture massive amounts of carbon dioxide and convert it into carbonate rock.

The robot armada, solar arrays, carbon-converting machines, and piles of rock would all grow exponentially, reaching "continental size in less than a decade," the paper concluded. Converting 20% of the carbon dioxide in the atmosphere would generate a layer of rock 50 centimeters (20 inches) thick covering a million square kilometers (390,000 square miles)—an area the size of Egypt.

The hitch, of course, is that self-replicating machines don't exist. Lackner moved on from that part of the plan, and briefly focused on solar power as a replacement for fossil fuels. But the more he studied the problem, the more he came to believe that renewable sources would struggle to compete with the price, abundance, and energy density of coal, oil, and gasoline.

"This suggested to me that fossil-fuel-based power will not just roll over and die," he says. But perhaps if carbon removal technologies were cheap enough, he thought, you could "force fossil-fuel providers to clean up after themselves."

A few years later, Lackner published a paper titled "Carbon Dioxide Extraction from Air: Is It an Option?" He argued that it was technically feasible and might be possible for as little as $15 a ton. (He now believes the price floor is probably between $30 and $50 a ton.)

In 2001 Lackner moved to Columbia University, where he cofounded Global Research Technologies, the first effort to commercialize direct air capture. Gary Comer, founder of the clothing and furniture company Lands' End, handed the company $8 million of what Lackner describes as "adventure capital, not venture capital."

The company built a small prototype but soon ran out of money. A group of investors bought the controlling interest, moved it to San Francisco, and renamed it Kilimanjaro Energy. Lackner served as an advisor and board member. But it quietly closed its doors after failing to raise more money.

Despite these failures, Lackner continued to try to figure out how to do air capture cheaply and efficiently. He's published more than 100 scientific papers and editorials on the subject, and applied for more than two dozen patents.

> SOME SCIENTIFIC CRITICS FOUND LACKNER'S PROJECTIONS NOT JUST WRONG BUT DANGEROUS. A PAIR OF CRITICAL PAPERS IN 2011 SOUNDED TO MANY LIKE A DEATH KNELL FOR DIRECT AIR CAPTURE. LACKNER WAS UNDAUNTED.

Some scientific critics, however, found Lackner's projections not just wrong but also dangerous. They feared that claiming direct air capture could be done cheaply and easily would reduce the pressure to slash emissions. In 2011, a pair of studies concluded that the technology would cost between $600 and $1,000 a ton.

Howard Herzog, a senior researcher at the MIT Energy Initiative, who coauthored one of the studies, took the added step of suggesting that "some purveyors" of the technology were "snake-oil salesmen." In an interview last year, Herzog told me he was mainly talking about Lackner. "He was the one who was really out there," he says.

Many read the two papers' conclusions as a death knell for direct air capture. Lackner stood firm, telling the journal Nature after the first of the studies was published: "They proved that one specific way to capture carbon dioxide from air is expensive. If you study penguins, you might jump to the conclusion that birds can't fly."

In 2014, he and his Global Research Technologies cofounder, Allen Wright, established the Center for Negative Carbon Emissions at Arizona State, where they've continued to try to get their own fledgling to take flight.

Planting synthetic forests

At the heart of the Center for Negative Carbon Emissions' design is a particular type of commercially available anion-exchange resin. As wind carries carbon dioxide in the air across those polymer strips, negatively charged ions bind with the gas molecules and convert them into bicarbonate—the main compound in baking soda and antacids.

The machine then retracts, pulling those saturated strips back into the container and pumping it full of water. The water begins converting the bicarbonate molecules into carbonate ions.

As the water drains away, those compounds become unstable and turn back into carbon dioxide in the air within the container. The now carbon dioxide–rich

GUTTER CREDIT HERE

The hundreds of polymer strips form a kind of sail that grabs carbon dioxide molecules as wind blows air through the device.

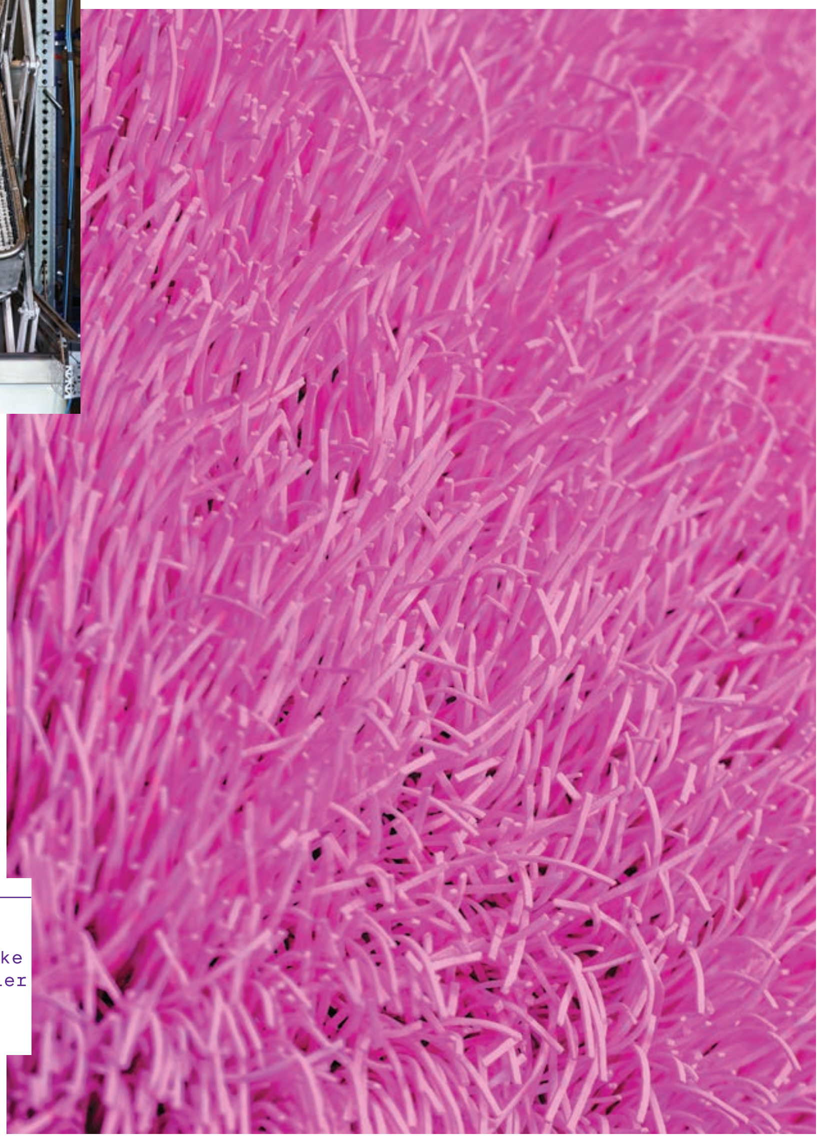

A close-up view of the carbon-capturing materials in a grass-like configuration, an earlier design that releases carbon dioxide when placed in a greenhouse.

"MY ARGUMENT HAS ALWAYS BEEN WE NEED TO BE PASSIVE," LACKNER SAYS. "WE WANT TO BE A TREE STANDING IN THE WIND AND HAVE THE CO_2 CARRIED TO US."

air can then be sucked out through a tube, and into an adjacent set of tanks.

Since carbon dioxide is relatively dilute in the air, most other direct capture approaches employ large fans to blow air over the binding materials to trap more of the gas. They then employ heat to drive the subsequent reactions that release the carbon dioxide. Both these steps use more energy. In contrast, Lackner says, his and Wright's approach just requires a little electricity to extend and retract the machine, pump the water, and vacuum out the air.

"My argument has always been we need to be passive," Lackner says. "We want to be a tree standing in the wind and have the CO_2 carried to us."

But there are big drawbacks to this method. It works only when the wind is blowing and makes sense only in dry areas, since humidity allows the carbon dioxide to escape. Moreover, the concentration of captured carbon in the resulting gas is less than 5%, compared with around 98% from a Carbon Engineering or Climeworks facility.

That low level is fine for fertilizing plants in greenhouses. But that's a tiny market, and Lackner has grander designs.

He envisions thousands of these machines plucking carbon dioxide from the sky in some dry and hot part of the world, while adjacent solar panels drive an electrolysis process that extracts hydrogen from water. The carbon dioxide and hydrogen could then be combined on site to produce thousands of barrels a day of synthetic fuel, which could be sold for heating or transportation, or used to feed the electric grid when renewables like wind and solar flag.

That plan, however, poses several challenges. Electrolysis is still very expensive. And they'd need to compress the carbon dioxide to the necessary concentration while removing water vapor, nitrogen, and oxygen.

That can be done, but it could substantially increase costs and energy needs. "This is a big, important piece that he's glossing over a bit," says Jennifer Wilcox, a professor at Worcester Polytechnic Institute and coauthor of the National Academies report.

Some believe Lackner's strengths as a theorist and big-picture guy haven't served him as well in translating those ideas into the necessary advances in materials science and chemistry. Notably, the Center for Negative Carbon Emissions project is trailing well behind Carbon Engineering, Climeworks, and Global Thermostat, which are amassing capital, hiring staffs, and building out demonstration if not commercial-scale facilities.

But Lackner remains confident that his approach will be less expensive than competing ones. "I can lay it out unit process by unit process, and in terms of first principles, at every step we're a little cheaper," he says.

Deep trouble

How does Lackner himself feel about the technology's prospects more than two decades after starting down this research path? It's not a simple answer. Lackner doesn't really do simple answers. During a walk across the university's palm-lined campus in Tempe, he says he remains confident that direct air capture is feasible and believes it could get much less expensive if it's able to reach commercial scale.

"But I'm less optimistic that we have the political will to go through that threshold," he says.

Given the high early costs and limited markets, he believes the technology will need significant government funding or tight regulations to be widely adopted—and more government support to cover the cost of capturing and burying the majority of the carbon dioxide that can't be used. He thinks we'll need to treat carbon dioxide like sewage, requiring consumers or companies to pay for its collection and disposal, whether in taxes or fees.

But after decades of relatively little political action on climate change, and fierce public resistance to carbon taxes, he fears the world isn't going to come around to that way of thinking until the suffering from climate catastrophes becomes too horrible to ignore.

What he is sure of, after spending more time than anyone else puzzling over carbon removal, is that we're going to need it. "I'm the first to admit that air capture isn't proven—and it certainly isn't proven at scale," Lackner says. "But we're in deep trouble if we can't figure it out."

James Temple is MIT Technology Review's senior editor covering climate change.

Your heart on your sleeve

By **DAN HON**

Being able to measure your heart's electrical activity at any point has revolutionary potential.

It begins seven years ago, when my doctor asks me whether I want to lose my foot. I say to him: *No, I do not want to lose my foot*. "Good," he says back: Monitor your blood sugar, keep it down, and we can manage this disease. Then *nobody* has to lose a foot.

It turns out I have type 2 diabetes, which—from a patient's point of view—boils down to a single data point: the amount of glucose in my bloodstream. Low is good; high is bad. Threatening my feet felt like a scare tactic, but the results of an undetected infection are very real for diabetics. We are often hit by a grim combination of weaker immune response and loss of feeling in the limbs, which can make a routine infection go very, very bad. And, like all 30 million Americans who have been diagnosed with diabetes, I face other potential complications, too: kidney, retinal, gum, and heart disease, never mind a high incidence of depression (unsurprisingly, it can be depressing to learn that you might lose a foot).

But yes, it's the foot that does it for me. That's when I start collecting health data.

I realize that for my entire life, I haven't paid much attention to my health. My body was just meat housing for my brain. Suddenly, with my FDA-approved glucose meter, I have a small device that tells me a number, and that number gives me a reason to care more about my body.

I begin to discover that it's not just glucose I can monitor. A range of data and devices can help me avoid other health problems. High blood pressure, for example, affects 75 million Americans and the majority of diabetics. I'm also at higher risk of AFib—atrial fibrillation, or an irregular heartbeat, which can increase the chance that I have a stroke.

Gathering this new information requires a patchwork of services, so I approach it like an engineer. I track steps using wearable devices from Fitbit and Nike, and apps like Moves. I watch for high blood pressure with a Withings smart monitor. The data is stored alongside my weight, body fat percentage, and body mass index, all measured with a smart scale. And all the time there's my blood glucose, measured six times a day, before and after each meal.

I export the data as CSVs and view it in hand-crafted graphs and dashboards. My ad hoc monitoring system makes me an early adopter, a bona fide member of the quantified-self movement.

Seven years later, though, my fringe obsession has become mainstream. My cobbled-together system has been replaced by Apple's shiny Health app, and I get prompted to exercise by a wearable that is more powerful than my first laptop. And my watch can even monitor my heart.

Photograph by Bruce Peterson

I've been wearing an Apple Watch for the last 15 months, using it to meet activity goals and monitor my health. ("Dan, you're so close to closing your Move ring. A brisk nine-minute walk should do it.") But the Series 4, Apple's latest model, has an extra function: a built-in electrocardiogram (ECG).

The gold-standard ECG measures the electrical activity of your heart with a 12-lead test, all wires and electrodes, administered by a medical professional. A watch that can run a basic version of this procedure—with a device you can wear all day, every day, for a price of a few hundred dollars—is a breakthrough.

Apple isn't the first to produce an over-the-counter ECG reader. AliveCor, a medical-device startup based in Silicon Valley, got there first with two FDA-licensed consumer ECG devices: the $100 KardiaMobile and the $199 Apple Watch band accessory KardiaBand.

All these devices are now used mainly to screen for AFib. That's a big deal, because not only do as many as 6.1 million Americans have the condition, but research suggests another 700,000 have irregular heartbeats that are undiagnosed. AFib contributes to an estimated 130,000 deaths each year in the US—but 20% of people whose strokes were due to AFib were unaware they had it until they were hospitalized. At the moment, even people with the best access to care get only two or three ECGs a year. Preventive screening could, if widely implemented, save thousands of lives.

Taking an ECG reading from a watch is a big step in that direction.

Not too much about the Series 4 feels different from the previous model— it's a little faster, and instead of a red dot on the digital crown, this one has a red circle. There's a rigorous tutorial that covers the notifications it can give me for irregular heart rhythms, and it takes me through the ECG app. Apple explains what an ECG is and, broadly, what it measures. It tells me the different results I might get, such as a normal heartbeat (known as sinus rhythm), AFib, and low or high heart rate. During setup, there are clean, easy-to-read screens telling me what the ECG can't do: detect a heart attack, blood clots, or other conditions like high blood pressure or high cholesterol. If I'm not feeling well, it says, I should talk to my doctor. If I'm experiencing chest pain, I should call emergency services. It's like the iTunes terms of service, but a lot shorter and *much* more serious.

Then it asks me to take a reading. This first time, I'm a little anxious. I remember that my mum has a history of hypertrophic cardiomyopathy, and my brother, too.

Apple Watch's ECG works by forming a circuit that runs from the back of the watch, where it touches the skin on my left wrist, through to the watch's crown, which I touch with a finger of my right hand. The app uses the electrical pulses running through this circuit to get my heart rate and, most important, to see if the upper and lower chambers of my heart are in rhythm. To take an ECG, I'll have to sit still and keep that right-hand finger on the digital crown for 30 seconds.

It's a long 30 seconds.

As the timer counts down, I feel the same anxiety mounting in my chest that I do when I have my blood pressure taken. I really want the upper and lower chambers of my heart to be in rhythm.

And then there it is, on my phone: "Setup Complete. This ECG does not show signs of atrial fibrillation."

I give an audible sigh of relief, and realize I've been holding my breath.

Over the next few weeks, I take my ECG a couple more times, but the urgency and anxiety have worn off. The only time I get a non-uniform result is when our family arrives at the airport at the start of our vacation. This one seems fine: I've had a stressful morning, and all the subsequent readings I take are back to normal.

In a month wearing the Series 4 that was loaned to me by Apple, the experience has been mostly mundane. That's probably how it'll go for most people. For a good friend of mine, though, the watch made a more dramatic difference.

It wasn't a surprise to hear that Tom had upgraded to the Series 4 when it came out. He's been an Apple user longer than I have, and he has a family history of AFib on his mother's side. (It turns out she already uses KardiaMobile, as well as hospital-style home monitoring.)

One day, while I was testing my own Apple Watch, Tom was deconstructing a rack of network equipment. He suddenly noticed his heart was pounding. Then he began feeling dizzy. Next came tunnel vision. He needed to sit down.

First he checked the pulse on his neck, but he realized his watch could provide more data. It said 203 beats per minute, so he fired up an ECG—the first time he'd done it, so he had to go through setup and onboarding first. When it took his reading, Tom's watch said it couldn't check for AFib because the heart rate was over 120 beats per minute: "If you're not feeling well, you should talk to your doctor," it said. Tom was definitely not feeling well, so he had a coworker take him to the hospital, where triage got him to a nurse straight away.

His nurse set up an ECG, the traditional "gold standard" kind, but Tom could feel that his heart rate had dropped closer to normal. He worried that the hospital test wouldn't find anything, so he unlocked his phone and passed the readings to the nurse, who showed them to the remote teledoctor on call.

"Oh, that's an SVT," the doctor said, immediately. A supraventricular tachycardia: an abnormally fast heartbeat caused by irregular electrical activity. The hospital ordered blood tests and sent Tom to his regular doctor for a follow-up.

This sequence of events encapsulates the promise of having a "good enough" ECG on demand: readings can be taken when symptoms happen, not after. The right data at the right time.

But Tom's experience feels fortuitous, too. What might have happened if Tom hadn't taken an ECG, or if there hadn't been a report for the doctor? Would the gold-standard hospital ECG have found anything?

The Withings Move ECG (above) and the KardiaMobile from AliveCor (below).

COURTESY PHOTOS

Those questions are moot. Tom *did* have an ECG, taken within seconds of his symptoms. He had more tests, and they showed he's got nothing to worry about for now. But he's been alerted to the danger. It worked. He's grateful.

Experience shows that when these devices are available, people use them. Fitbit devices now track more than 25 million active users. In early 2019, the connected-device maker Withings announced that its forthcoming watch will have an ECG reader. Apple alone sells millions of watches each year. Consumer ECGs are here, and they're probably going to get cheaper and more ubiquitous.

These systems are creating a mountain of health data, though. How do we interpret this information? Can the medical profession cope with the volume? There is no excess of experienced cardiologists waiting to review the 20 million ECGs AliveCor recorded in 2017, and that was before Apple turned up.

It seems inevitable that we'll throw deep-learning algorithms at the data and look for new ways to use it. Apple recently announced a study with Johnson & Johnson to screen for stroke risk. And AliveCor's KardiaK software—developed through a partnership with the Mayo Clinic—has been granted accelerated clearance by the FDA. KardiaK uses deep learning on ECGs to screen for hyperkalemia, or elevated potassium levels in the blood. For people with kidney disease, the condition comes with a higher risk of arrhythmia and death.

For all the potential benefit, though, one could envision things getting quickly out of hand. In a few product cycles, anything from a $25 Xiaomi wearable to a high-end Apple Watch could be collecting a range of health information and using it to screen for conditions like hypertension, sleep apnea, diabetes, or even changes in mood. In a tired joke, I imagine a future continuously monitored by Microsoft's Clippy: "It looks like you're starting to get depressed. Would you like help getting some exercise?"

How prepared are we to deal with the ethical issues these predictive models create? How can the technologies be audited to make sure they work for all users and not, accidentally, just for subsets of populations? When we use this data—and it is when, not if—we need to be able to answer these questions, and others.

Seven years ago I started tracking my blood sugar because I didn't want to lose a foot. Now, after a month of using the Series 4 Apple Watch, I'm reminded what data can mean for my heart and, by extension, my mind.

The red dot on the digital crown of my Series 3 Watch was comforting. It meant that I had cell coverage and wasn't out of touch. Now, the red circle on the Series 4 feels even more reassuring—but in an entirely different way.

Dan Hon is a product strategist working on California's digital services and an occasional technology writer based in Portland, Oregon.

By **NIALL FIRTH**

Illustration by Dingding Hu

The meat without the cow

Meat production spews tons of greenhouse gas and uses up too much land and water. Is there an alternative that won't make us do without?

Memphis Meats CEO Ulma Valeti (center) and chief science officer Nicholas Genovese (right) watch a chef prepare one of their creations.

MEMPHIS MEATS

In 2013, the world's first burger from a lab was cooked in butter and eaten at a glitzy press conference. The burger cost £215,000 ($330,000 at the time) to make, and despite all the media razzmatazz, the tasters were polite but not overly impressed. "Close to meat, but not that juicy," said one food critic.

Still, that one burger, paid for by Google cofounder Sergey Brin, was the earliest use of a technique called cellular agriculture to make edible meat products from scratch—no dead animals required. Cellular agriculture, whose products are known as cultured or lab-grown meat, builds up muscle tissue from a handful of cells taken from an animal. These cells are then nurtured on a scaffold in a bioreactor and fed with a special nutrient broth.

A little over five years later, startups around the world are racing to produce lab-grown meat that tastes as good as the traditional kind and costs about as much.

They're already playing catch-up: "plant-based" meat, made of a mix of non-animal products that mimic the taste and texture of real meat, is already on the market. The biggest name in this area: Impossible Foods, whose faux meat sells in more than 5,000 restaurants and fast food chains in the US and Asia and should be in supermarkets later this year. Impossible's research team of more than 100 scientists and engineers uses techniques such as gas chromatography and mass spectrometry to identify the volatile molecules released when meat is cooked.

The key to their particular formula is the oxygen-carrying molecule heme, which contains iron that gives meat its color and metallic tang. Instead of using meat, Impossible uses genetically modified yeast to make a version of heme that is found in the roots of certain plants.

Impossible has a few competitors, particularly Beyond Meat, which uses pea protein (among other ingredients) to replicate ground beef. Its product is sold in supermarket chains like Tesco in the UK and Whole Foods in the US, alongside real meat and chicken. Both Impossible and Beyond released new, improved versions of their burgers in mid-January.

In contrast, none of the lab-grown-meat startups has yet announced a launch date for its first commercial product. But when that happens—some claim as early as the end of this year—the lab-grown approach could turn the traditional meat industry on its head.

"I suspect that cultured meat proteins can do things that plant-based proteins can't in terms of flavor, nutrition, and performance," says Isha Datar, who leads New Harvest, an organization that helps fund research in cellular agriculture. Datar, a cell biologist and a fellow at the MIT Media Lab, believes cultured meats will more closely resemble real meat, nutritionally and functionally, than the plant-based kinds do. The idea is that a die-hard carnivore (like me) might not feel so put off at the thought of giving up the real thing.

A GLOBAL RISK

You might ask, why would anyone want to? The answer is that our meat consumption habits are,

in a very literal sense, not sustainable. Livestock raised for food already contribute about 15% of the world's global greenhouse-gas emissions. (You may have heard that if cows were a country, it would be the world's third biggest emitter.) A quarter of the planet's ice-free land is used to graze them, and a third of all cropland is used to grow food for them. A growing population will make things worse. It's estimated that with the population expected to rise to 10 billion, humans will eat 70% more meat by 2050. Greenhouse gases from food production will rise by as much as 92%.

In January a commission of 37 scientists reported in The Lancet that meat's damaging effects not only on the environment but also on our health make it "a global risk to people and the planet." In October 2018 a study in Nature found that we will need to change our diets significantly if we're not to irreparably wreck our planet's natural resources.

We'll need to change our diets to avoid wrecking the planet.

"Without changes toward more plant-based diets," says Marco Springmann, a researcher in environmental sustainability at the University of Oxford and the lead author of the Nature paper, "there is little chance to avoid dangerous levels of climate change."

The good news is that a growing number of people now seem to be rethinking what they eat. A recent report from Nielsen found that sales of plant-based foods intended to replace animal products were up 20% in 2018 compared with a year earlier. Veganism, which eschews not just meat but products that come from greenhouse-gas-emitting dairy livestock too, is now considered relatively mainstream.

That doesn't necessarily equate to more vegans. A recent Gallup poll found that the number of people in the US who say they are vegan has barely changed since 2012 and stands at around just 3%. Regardless, Americans are eating less meat, even if they're not cutting it out altogether.

AND NOW FOR THE LAWSUITS

Investors are betting big that this momentum will continue. Startups such as MosaMeat (cofounded by Mark Post, the scientist behind the £215,000 burger), Memphis Meats, Supermeat, Just, and Finless Foods have all swept up healthy sums of venture capital. The race now is to be first to market with a palatable product at an acceptable cost.

Memphis Meats' VP of product and regulation, Eric Schulze, sees his product as complementing the real-meat industry. "In our rich cultural tapestry as a species, we are providing a new innovation to weave into our growing list of sustainable food traditions," he says. "We see ourselves as an 'and,' not 'or,' solution to helping feed a growing world."

The traditional meat industry doesn't see it that way. The National Cattlemen's Beef Association in the US dismissively dubs these new approaches "fake meat." In August 2018, Missouri enacted a law

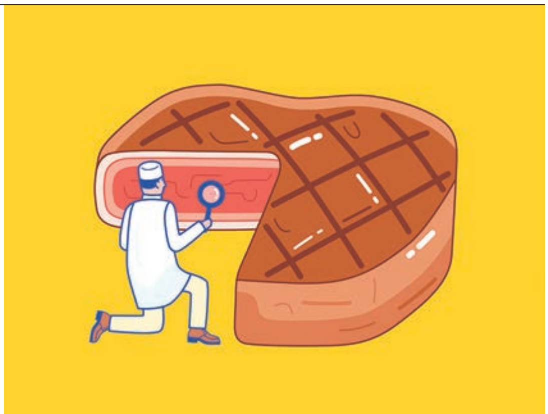

that bans labeling any such alternative products as meat. Only food that has been "derived from harvested production of livestock or poultry" can have the word "meat" on the label in any form. Breaking that law could lead to a fine or even a year's jail time.

The alternative-meat industry is fighting back. The Good Food Institute, which campaigns for regulations that favor plant-based and lab-grown meats, has joined forces with Tofurky (the makers of a tofu-based meat replacement since the 1980s), the American Civil Liberties Union, and the Animal Legal Defense Fund to get the law overturned. Jessica Almy, the institute's policy director, says the law as it stands is "nonsensical" and an "affront" to the principle of free speech. "The thinking behind the law is to make plant-based meat less appealing and to disadvantage cultured meat when it comes on the market," she says.

Almy says she's confident their case will be successful and is expecting a temporary injunction to be granted soon. But the Missouri battle is just the start of a struggle that could last years. In February 2018, the US Cattlemen's Association launched a petition that calls on the US Department of Agriculture (USDA) to enact a similar federal law.

Traditional meat-industry groups have also been very vocal on how cultured meat and plant-based meats are to be regulated. Last summer a group of the biggest agricultural organizations in the US (nicknamed "The Barnyard") wrote to President Trump asking for reassurance that the USDA will oversee cultured meat to ensure "a level playing field." (The USDA has tougher, more stringent safety inspections than the Food and Drug Administration.)

In November 2018, the USDA and the FDA finally released a joint statement to announce that the two regulators would share the responsibilities for overseeing lab-grown meats.

"I think there will be lines outside the store that are longer than for the next iPhone."

THE BOVINE SERUM PROBLEM

Some cultured-meat startups say this confusion over regulations is the only thing holding them back. One firm, Just, says it plans to launch a ground "chicken" product this year and has trumpeted a partnership with a Japanese livestock firm to produce a "Wagyu beef" product made from cells in the lab. Its CEO is Josh Tetrick, who'd previously founded the controversial startup Hampton Creek, Just's forebear. (The FDA had at one time banned the firm from calling its signature product mayonnaise, as it did not contain any eggs.) Speak to Tetrick, a bullish, confident young man, and you get a sense of the drive and excitement behind the alternative-meat market. "The only [limit] to launching," he says, "is regulatory."

That's optimistic, to say the least. The lab-meat movement still faces big technical hurdles. One is that making the product requires something called fetal bovine serum. FBS is harvested from fetuses taken from pregnant cows during slaughter. That's an obvious problem for a purportedly cruelty-free product. FBS also happens to be eye-wateringly expensive. It is used in the biopharmaceutical industry and in basic cellular research, but only in tiny amounts. Cultured meat, however, requires vast quantities. All the lab-meat startups will have to use less of it—or eliminate it completely—to make their products cheap enough. Last year Finless Foods (which aims to make a fish-free version of bluefin tuna) reported that it had halved the amount of FBS it needs to grow its cells. And Schulze says the Memphis Meats team is working on ways of cutting it out entirely.

But there are other issues, says Datar, of New Harvest. She says we still don't understand the fundamental processes well enough. While we have quite a deep understanding of animals used in medical research, such as lab mice, our knowledge of agricultural animals at a cellular level is rather thin. "I'm seeing a lot of excitement and VCs investing but not seeing a lot in scientific, material advancements," she says. It's going to be tricky to scale up the technology if we're still learning how these complex biological systems react and grow.

Lab-grown meat has another—more tangible—problem. Growing muscle cells from scratch creates pure meat tissue, but the result lacks a vital component of any burger or steak: fat. Fat is what gives meat its flavor and moisture, and its texture is hard to replicate. Plant-based meats are already getting around the problem—to some extent—by using shear cell technology that forces the plant protein mixture into layers to produce a fibrous meat-like texture. But if you want to create a meat-free "steak" from scratch, some more work needs to be done. Cultured meat will need a way to grow fat cells and somehow mesh them with the muscle cells for the end result to be palatable.

That has proved tricky so far, which is the main reason that first burger was so mouth-puckeringly dry.

The scientists at the Netherlands-based cultured-meat startup Meatable might have found a way. The team has piggybacked on medical stem-cell research to find a way of isolating pluripotent stem cells in cows by taking them from the blood in umbilical cords of newborn calves. Pluripotent cells, formed early in an embryo's development, have the ability to develop into any type of cell in the body. This means they can also be coaxed into forming fat, muscle, or even liver cells in lab-grown meat.

Meatable's work might mean that the cells can be tweaked to produce a steak-like product whose fat and muscle content depends on what the customer prefers: a rib-eye steak's characteristic marbling, for example. "We can add more fat, or make it leaner—we can do anything we want to. We have new control over how we feed the cells," says Meatable CTO Daan Luining, who is also a research director at the nonprofit Cellular Agriculture Society. "Pluripotent cells are like the hardware. The software you're running turns it into the cell you want. It's already in the cell—you just need to trigger it."

But the researchers' work is also interesting because they have found a way to get around the FBS problem: the pluripotent cells don't require the serum to grow. Luining is clearly proud of this. "To circumvent that using a different cell type was a very elegant solution," he says.

He concedes that Meatable is still years away from launching a commercial product, but he's confident about its eventual prospects. "I think there will be lines outside the store that are longer than for the next iPhone," he says.

IF YOU MAKE IT, WILL THEY EAT IT?

As it stands, lab-grown meat is not quite as virtuous as you might think. While its greenhouse emissions are below those associated with the biggest villain, beef, it is more polluting than chicken or the plant-based alternatives, because of the energy currently required to produce it. A World Economic Forum white paper on the impact of alternative meats found that lab-grown meat as it is made now would produce only about 7% less in greenhouse-gas emissions than beef. Other replacements, such as tofu or plants, produced reductions of up to 25%. "We will have to see if companies will really be able to offer low-emissions products at reasonable costs," says Oxford's Marco Springmann, one of the paper's coauthors.

What your food does to the planet

Kilograms of carbon dioxide equivalent* per 200 calories

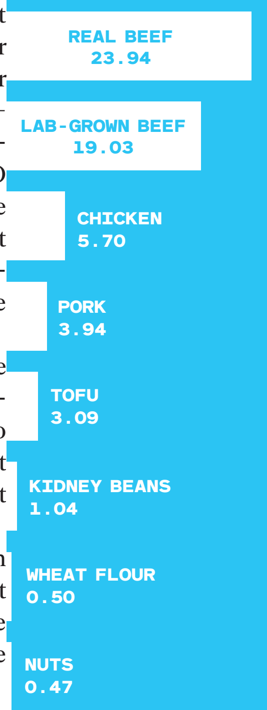

*A CO_2 equivalent is a metric that allows different types of greenhouse gases to be measured on the same scale.

Source: World Economic Forum

It is also unclear how much better for you lab-grown meat would be than the real thing. One reason meat has been linked to a heightened cancer risk is that it contains heme, which could also be present in cultured meats.

And will people even want to eat it? Datar thinks so. The little research there has been on the subject backs that up. A 2017 study published in the journal PLoS One found that most consumers in the US would be willing to try lab-grown meat, and around a third were probably or definitely willing to eat it regularly.

Expecting the whole world to go vegan is unrealistic. But a report in Nature in October 2018 suggested that if everyone moved to the flexitarian lifestyle (eating mostly vegetarian but with a little poultry and fish and no more than one portion of red meat a week), we could halve the greenhouse-gas emissions from food production and also reduce other harmful effects of the meat industry, such as the overuse of fertilizers and the waste of fresh water and land. (It could also reduce premature mortality by about 20%, according to a study in The Lancet in October, thanks to fewer deaths from ailments such as coronary heart disease, stroke, and cancer.)

Some of the biggest players in the traditional meat industry recognize this and are subtly rebranding themselves as "protein producers" rather than meat companies. Like Big Tobacco firms buying vape startups, the meat giants are also buying stakes in this new industry. In 2016, Tyson Foods, the world's second biggest meat processor, launched a venture capital fund to support alternative-meat producers; it's also an investor in Beyond Meat. In 2017, the third biggest, Cargill, invested in cultured-meat startup Memphis Meats, and Tyson followed suit in 2018. Many other big food producers are doing the same; in December 2018, for example, Unilever bought a Dutch firm called the Vegetarian Butcher that makes a variety of non-meat products, including plant-based meat substitutes.

"A meat company doesn't do what they do because they want to degrade the environment and don't like animals," says Tetrick, the Just CEO. "They do it because they think it's the most efficient way. But if you give them a different way to grow the company that's more efficient, they'll do it."

At least some in the meat industry agree. In a profile last year for Bloomberg, Tom Hayes, then the CEO of Tyson, made it clear where he saw the company's eventual future. "If we can grow the meat without the animal," he said, "why wouldn't we?"

Niall Firth is MIT Technology Review's news editor.

By **LEIGH PHILLIPS**

It's time to reconsider the new

nuclear option

Facing up to our climate challenge may require a fresh generation of nuclear power. Luckily, advances are on the horizon.

Photograph by Julian Berman

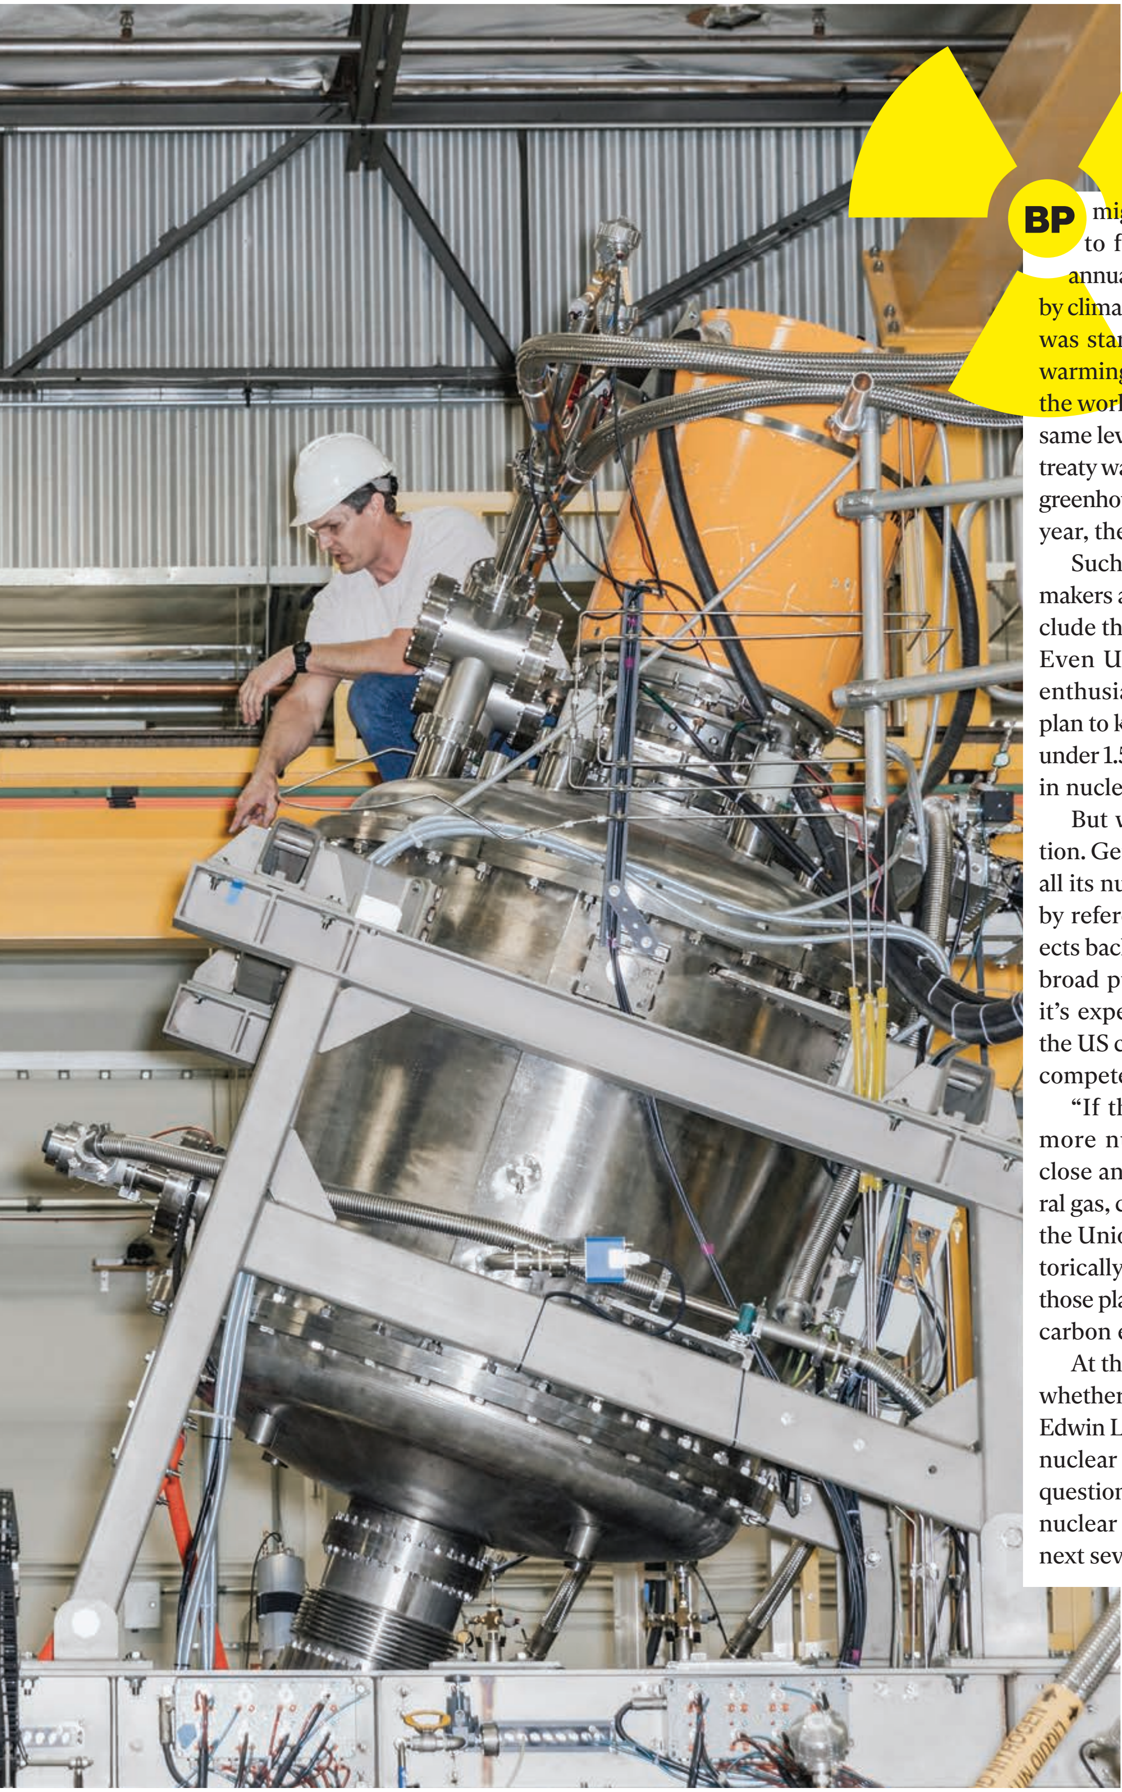

BP might not be the first source you go to for environmental news, but its annual energy review is highly regarded by climate watchers. And its 2018 message was stark: despite the angst over global warming, coal was responsible for 38% of the world's power in 2017—precisely the same level as when the first global climate treaty was signed 20 years ago. Worse still, greenhouse-gas emissions rose by 2.7% last year, the largest increase in seven years.

Such stagnation has led many policymakers and environmental groups to conclude that we need more nuclear energy. Even United Nations researchers, not enthusiastic in the past, now say every plan to keep the planet's temperature rise under 1.5 °C will rely on a substantial jump in nuclear energy.

But we're headed in the other direction. Germany is scheduled to shut down all its nuclear plants by 2022; Italy voted by referendum to block any future projects back in 2011. And even if nuclear had broad public support (which it doesn't), it's expensive: several nuclear plants in the US closed recently because they can't compete with cheap shale gas.

"If the current situation continues, more nuclear power plants will likely close and be replaced primarily by natural gas, causing emissions to rise," argued the Union of Concerned Scientists—historically nuclear skeptics—in 2018. If all those plants shut down, estimates suggest, carbon emissions would increase by 6%.

At this point, the critical debate is not whether to support existing systems, says Edwin Lyman, acting director of the UCS's nuclear safety project. "A more practical question is whether it is realistic that new nuclear plants can be deployed over the next several decades at the pace needed."

A photograph taken in 2016 shows the central confinement vessel of a prototype fusion reactor built by Tri Alpha Energy (now TAE Technologies).

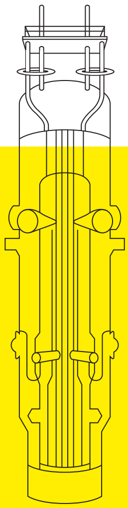

Small modular reactors

SMRs are a slimmed-down version of conventional fission reactors. Although they produce far less power, their smaller size and use of off-the-shelf components help reduce cost.

COMPANIES	NuScale Power
POWER OUTPUT	50-200 megawatts
EXPECTED LIFE SPAN	60 years
COST	$100 million prototype, $2 billion to develop
AVAILABLE	2026

As of early 2018 there were 75 separate advanced fission projects trying to answer that question in North America alone, according to the think tank Third Way. These projects employ the same type of reaction used in the conventional nuclear reactors that have been used for decades—fission, or splitting atoms.

One of the leading technologies is the small modular reactor, or SMR: a slimmed-down version of conventional fission systems that promises to be cheaper and safer. NuScale Power, based in Portland, Oregon, has a 60-megawatt design that's close to being deployed. (A typical high-cost conventional fission plant might produce around 1,000 MW of power.)

NuScale has a deal to install 12 small reactors to supply energy to a coalition of 46 utilities across the western US, but the project can go ahead only if the group's members agree to finance it by the end of this year. History suggests that won't be easy. In 2011, Generation mPower, another SMR developer, had a deal to construct up to six reactors similar to NuScale's. It had the backing of corporate owners Babcock & Wilcox, one of the world's largest energy builders, but the pact was shelved after less than three years because no new customers had emerged. No orders meant prices wouldn't come down, which made the deal unsustainable.

While NuScale's approach takes traditional light-water-cooled nuclear reactors and shrinks them, so-called generation IV systems use alternative coolants. China is building a large scale sodium-cooled reactor in Fujian province that's expected to begin operation by 2023, and Washington-based TerraPower has been developing a sodium-cooled system that can be powered with spent fuel, depleted uranium, or uranium straight out of the ground. TerraPower—Bill Gates is an investor—forged an agreement with Beijing to construct a demonstration plant by 2022, but the Trump administration's restrictions on Chinese trade make its future questionable.

Another generation IV variant, the molten-salt reactor, is safer than earlier designs because it can cool itself even if the system loses power completely. Canadian company Terrestrial Energy plans to build a 190 MW plant in Ontario, with its first reactors producing power before 2030 at a cost it says can compete with natural gas.

One generation IV reactor could go into operation soon. Helium-cooled, very-high-temperature reactors can run at up to 1,000 °C, and the state-owned China National Nuclear Corporation has a 210 MW prototype in the eastern Shandong province set to be connected to the grid this year.

For many, though, the great energy hope remains nuclear fusion. Fusion reactors mimic the nuclear process inside the sun, smashing lighter atoms together to turn them into heavier ones and releasing vast amounts of energy along the way. In the sun, that process is powered by gravity. On Earth, engineers aim to replicate fusion conditions with unfathomably high temperatures—on the order of 150 million °C—but they have found it hard to confine the plasma required to fuse atoms.

One solution is being built by ITER, previously known as the International Thermonuclear Experimental Reactor, under construction since 2010 in Cadarache, France. Its magnetic confinement system has global support, but costs have exploded to $22 billion amid delays and political wrangling. The first experiments, originally scheduled for 2018, have been pushed back to 2025.

Vancouver's General Fusion uses a combination of physical pressure and magnetic fields to create plasma pulses that last millionths of a second. This is a less complicated approach than ITER's, making it far cheaper—but technical challenges remain, including making titanium components that can handle the workload. Still, General Fusion expects its reactors to be deployable in 10 to 15 years.

California-based TAE Technologies, meanwhile, has spent 20 years developing a fusion reactor that converts energy directly into electricity. The company, which has received $500 million from investors, predicted in

KYLE THOMAS HEMINGWAY

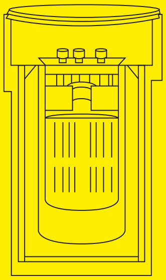

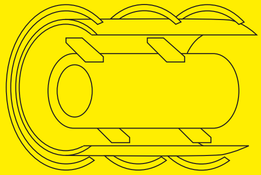

Advanced fission	Fusion
These reactors are designed to be safer than traditional water-cooled reactors, using coolants such as liquid sodium or molten salts instead. Most advanced is the "pebble bed" reactor, cooled by a gas such as helium; China is ready to connect the first such reactor to the grid this year.	Technical progress is still slow after decades of investment, but fusion companies are focused on how to contain the plasma required to replicate the thermonuclear conditions of the sun. Techniques include magnetic confinement, which traps plasma continuously at low pressure; inertial confinement, using lasers and pulsing plasma for nanoseconds at a time; and magnetized target fusion, which combines the two with pulses of plasma controlled by magnets.
China National Nuclear Corporation, TerraPower, Terrestrial Energy	ITER, TAE Technologies, General Fusion, Commonwealth Fusion Systems
190-600 megawatts	100-500 megawatts
40-60 years	35 years
Pebble beds: $400 million to $1.2 billion Sodium-cooled and molten salt: $1 billion prototype	ITER: currently $22 billion Cost of a commercial version is unknown
Pebble bed in 2019; sodium-cooled 2025; molten salt 2030	No earlier than 2035

January that it would be commercial within five years.

So will any of these technologies succeed? Advanced fission reduces nuclear waste—even using it as fuel—and drastically shrinks the chance of tragedies like Fukushima or Chernobyl. Yet no such reactors have been licensed or deployed outside China or Russia. Many voters simply don't believe companies when they promise that new technologies can avoid old mistakes.

It's not just politics, though: cost is also a factor. Advanced fission promises to reduce the ridiculously expensive up-front costs of nuclear energy by creating reactors that can be factory built, rather than custom made. This would cause prices to plummet, just as they have for wind and solar. But private companies have rarely proved successful at bringing these projects to completion: the biggest advances have come from highly centralized, state-driven schemes that can absorb risk more easily.

General Fusion CEO Chris Mowry argues that fission simply faces too many barriers to be successful. He has experience: he was a founder of mPower, the SMR company that was mothballed in 2014. Fusion reactors might be harder to build, he suggests, but they are more socially acceptable. This is why there's been a rush of venture capital into fusion, he says—investors are confident there will be a sea of eager buyers waiting for whoever can make it work first.

But does fusion really have that much more room to maneuver? It's true that the low-level, short-lived radioactive tritium waste it produces represents no serious danger, and the technology means that meltdowns are impossible. But costs are still high and time lines are still long—ITER's fusion reactor is massively more expensive than originally planned and won't be workable for at least 15 years. Meanwhile, Green politicians in Europe already want ITER shut down, and many anti-nuclear campaigners don't distinguish between fission and fusion.

Experts might be lining up behind nuclear, but convincing skeptical voters is something else.

Many voters simply don't believe companies' promises that new technologies can avoid old mistakes.

Leigh Phillips is a science writer based in British Columbia, Canada.

By **BONNIE ROCHMAN**

The search for a simple preemie predictor

Complications from preterm birth are the leading cause of death worldwide in children under five.

Fifteen million babies are born prematurely each year. Stephen Quake's daughter, Zoe, was one of them: she arrived via emergency C-section after Quake and his wife, Athina, made a middle-of-the-night dash to the emergency room, a month before Zoe was due. She spent her first night in an incubator, and her father, a bioengineer then at Caltech, wondered why birth couldn't be more predictable.

That question lingered in Quake's mind. Months before Zoe began her junior year of high school, her dad announced he had developed a maternal blood test that may be able to alert women that they are going to deliver prematurely—before 37 completed weeks of gestation. He has since launched a startup to commercialize the technology and create a cheap, easy test that women could take around the sixth month of pregnancy.

The prematurity test isn't Quake's first foray into prenatal health. When Athina was pregnant with Zoe, she had undergone amniocentesis, an invasive needle biopsy used to detect Down syndrome and other conditions. When it's executed by doctors with lots of experience, the risk of miscarriage is low, but it exists—and that's nerve-racking for expectant parents. "I thought, Oh my God, this is awful—that you have to risk losing the baby to ask a diagnostic question," he says.

Convinced there had to be a better way, Quake got to work developing noninvasive blood tests to assess much of the same information as amniocentesis but with less risk to the pregnancy. He used bits of free-floating fetal DNA found in maternal blood to get a peek at the genetic makeup of the fetus. More than a decade later, multiple biotech companies offer a version of similar tests for Down syndrome and other conditions to pregnant women in clinics worldwide.

Likewise, blood tests, often called "liquid biopsies," are in development for a number of applications, including detecting early-stage cancer and revealing whether a replacement heart is failing in the body of a transplant recipient. In 2014, Quake identified evidence of dying neurons in the blood circulation of Alzheimer's patients, a step that is being used to develop tests for neurodegenerative and autoimmune diseases.

Predicting preterm birth would be another important breakthrough. Globally, more than one in 10 babies is born preterm, a public health problem that cuts across socioeconomic and geographic boundaries. Babies in poor nations like Malawi are born too soon—the country has an 18% rate of preterm birth, the highest in the world—but so are babies in the US, like Quake's daughter in prosperous Southern California.

Complications from preterm birth are the leading cause of death worldwide in children under the age of five. Preterm babies can struggle with infection, learning disabilities, and problems with vision and hearing. In poor countries, babies born significantly preterm often don't survive. In wealthy countries they usually do, but sometimes with long-term consequences including behavioral problems and neurological disorders such as cerebral palsy. There's an economic factor, too: babies born preterm cost, on average, 10 times as much over the first year of life as those whose birth had no complications.

Just ask Jen Sinconis, whose twins arrived with no warning at 24 weeks' gestation in 2006. Twin pregnancies are considered high risk, but Sinconis's pregnancy had been uneventful until she started having what she assumed were Braxton Hicks contractions, which can occur weeks in advance of delivery as the uterus primes itself for labor. She was wrong, and her twin boys arrived within six hours. Aidan

"I thought, Oh my God, this is awful—that you have to risk losing the baby to ask a diagnostic question."


GETTY IMAGES

Jen Sinconis's twins arrived at 24 weeks in 2006. Now 12, the boys are mostly healthy. Above, one of the boys in the ICU.

weighed 1 pound, 14 ounces (850 grams) and had to spend three months in the hospital; Ethan weighed 1 pound, 6 ounces, and was worse off. He was on oxygen for most of his first year of life and barely escaped needing a tracheotomy. Sinconis received a shot of surfactant to help develop her sons' lungs as soon as she reached the hospital, but if a test had been able to alert her doctor that she was at risk for early labor, she could have been given the medicine sooner, when it could possibly have made a difference. "If I had known they would have been born prematurely, our entire life would be different," says Sinconis, a creative producer at Starbucks corporate headquarters in Seattle.

The boys' medical care cost more than $2 million and didn't end when they left the hospital. They remained in isolation at home for the first three and a half years of their lives; Sinconis can barely keep track of the number of doctors and therapists they've seen through the years. She and her husband were forced to sell their home, liquidate their retirement and savings accounts, and eventually declare bankruptcy to deal with the nearly $450,000 that insurance wouldn't cover. Now 12, the boys have mostly caught up developmentally to other children their age. But their parents are just starting to emerge from their financial struggles. "We're way overdue for a way to predict preterm birth," Sinconis says.

A NEW TEST

Zoe, now 17, "is all grown up and totally healthy," says Quake, a professor at Stanford University for the past 14 years, but figuring out how to predict preterm birth had been in the back of his mind since she was born. It "felt like the next big mountain to climb," he says. "We had gained confidence from noninvasive prenatal testing. Preterm birth was like Mt. Everest."

Quake knew there were no meaningful diagnostics that could identify which pregnant women would give birth too soon. The biggest tip-off is having given birth to a preterm baby before, something of little use for a first-time mom. Additionally, preterm delivery can be caused by multiple factors: infection, twins, or even maternal stress.

"We don't have any understanding about what is triggering preterm birth," says Ronald Wapner, director of reproductive genetics at Columbia University Irving Medical Center. "We have been shotgunning it."

Quake also knew that direct DNA measurements wouldn't help. Analyzing a baby's DNA, inherited from his or her parents, is fundamental to testing for Down syndrome because it can reveal the presence of an extra chromosome. "It's a genetic question," says Quake. But research has shown that the baby's genetic profile makes a minimal contribution to prematurity. So instead, Quake focused on DNA's molecular cousin, RNA. These molecules are harder to spot in blood (they're short-lived) but would provide a more relevant readout, Quake believed, because their levels go up and down according to what's going on in a person's body. Could it be that a pregnancy headed for trouble was sounding early alarm signals?

Quake and his team, including Mira Moufarrej, a grad student in his lab, scrutinized blood samples from 38 African-American women considered at risk for preterm birth, in some cases because they'd previously had a premature baby. Overall, black children in the US are born prematurely about 50% more often than whites. Thirteen of the women ended up delivering early. By analyzing

FAMILY PHOTOS COURTESY OF JENNIFER SINCONIS

RNA molecules in their blood, the researchers found seven genes whose changing activity signals, taken together, seemed to predict which babies had arrived prematurely.

Quake told me he was surprised by the result. "Holy shit, might we have figured out a way to determine preterm birth?" he recalls thinking. "We're still trying to understand the biology behind these seven genes," he adds; it's not yet clear whether the signals are emanating from the mother, the placenta, or the baby. Quake suspects they are "reflecting the mom's response to the pregnancy going off track." In other words, he says, "the whole thing is derailing and the mom is responding to that."

> "Holy shit, might we have figured out a way to determine preterm birth?" he recalls thinking.

"The beauty of this approach is that it allows us to see a conversation going on between the mother, the fetus, and the placenta," says David Stevenson, co-director of Stanford's Maternal and Child Health Research Institute and principal investigator at its prematurity research center. "It's like eavesdropping. Now we can access this as it's being communicated, which helps us understand what's going on throughout pregnancy."

TREATMENT HOPE

Five hundred years ago, fascinated by his anatomical dissection of the womb of a pregnant women who had died, Leonardo da Vinci wrote about his intention to unravel the secrets behind conception and preterm birth. He never did, and even today, there are relatively few answers. Perhaps because so little is known, pharmaceutical companies haven't seen preterm birth as a promising area for investment. Indeed, it is "one of the most neglected issues," says Sindura Ganapathi, co-leader of the Maternal, Newborn & Child Health Discovery & Tools portfolio at the Gates Foundation, which along with the March of Dimes and the CZ Biohub, a medical initiative funded by Mark Zuckerberg and his wife Priscilla Chan, has funded Quake's work.

"We need many more interventions," says Ganapathi. "We are pretty limited in our armamentarium."

A test could be a first step toward new drugs or treatments. Knowing who is at risk would let women prepare—say, by picking a hospital with a neonatal intensive care unit or working with an obstetrician who could prescribe progesterone, a drug sometimes given to try to extend pregnancy. "It goes back to personalized treatment," says Wapner. "We still haven't been able to identify how progesterone works and who it works for better. RNA could help us better understand who should get these medications."

The new window on pregnancy could lead to applications beyond preterm birth. "From the standpoint of where this could go, you could look at placental development, fetal development, and fetal-maternal interaction," says Wapner. "RNA has been the stepsister of DNA until very recently. It's a damn good clue about how to differentiate who's at risk of preterm birth, and it could give us a better way of evaluating what's going on during pregnancy."

In line with that, Quake has formed a startup, called Akna Dx, with lofty goals. It's raised more than $10 million from investors including Khosla Ventures of Menlo Park, California. "Our idea is to do blood-based tests to give key insights," says CEO and cofounder Maneesh Jain. "What is a fetus's gestational age? Are you at risk for preterm birth, or severe postpartum depression? Pregnancy tends to still be a big black box. We want to give you insights into what is happening internally so you can take action."

Other experts say more evidence is needed that RNA can provide those insights. That's because so many different factors can contribute to prematurity, and it's not clear how well Quake's biomarkers will do in a broader population. "The difficulty is that preterm delivery is not caused by one thing," says Diana Bianchi, director of the Eunice Kennedy Shriver National Institute of Child Health and Human Development and an expert in noninvasive prenatal testing. Infection, a compromised placenta, maternal stress, a twin pregnancy—all of these and more can trigger preterm birth. "In really small numbers, Steve was accurately able to distinguish women at risk of delivering preterm," says Bianchi. "But the numbers were really small."

Quake readily agrees that his initial findings need to be validated through a large clinical trial before any test would be ready for commercial use. Quake's team is working to confirm that the results from the African-American women hold up in other groups as well. Collaborators, including some of Akna's cofounders, are now collecting blood samples from 1,000 pregnant women.

"We hope this is going to save a lot of lives," says Quake. "That's really what we're aiming for. But this is just the beginning of the story … It's a very fertile area, no pun intended."

Bonnie Rochman is a health and science writer based in Seattle and the author of The Gene Machine.

We sat down to talk about breakthrough technologies, China, and reasons to be cheerful with this issue's guest editor.

By **Gideon Lichfield**
Photograph by Ian Allen

Bill Gates explains why we should all be optimists

These are edited excerpts from a conversation with Gates at his Seattle office on January 9. You can watch the full interview at **technologyreview.com/billgates**.

There are a lot of worries today about technology's harmful effects. How do you retain your famous optimism about it?
Look at how long people are living, the reduction of under-five mortality, the reduction in how poorly women are treated. Globally, inequity is down: poorer countries are getting richer faster than the richer countries are getting richer. The bulk of humanity lives in middle-income countries today. Fifty years ago, there were very few middle-income countries. Then there's the ability of science to solve problems. In heart disease and cancer we've made a lot of progress; in some of the more chronic diseases like depression and diabetes … Even in obesity, we're gaining some fundamental understandings of the microbiome and the signaling mechanisms involved.

So, yes, I am optimistic. It does bother me that most people aren't.

Maybe you have successful person's bias?
Of course, we have to factor that in. In my own life I've been extremely lucky. But even subtracting out my personal experience, I think the big picture is that it's better to be born today than ever, and it will be better to be born 20 years from now than today.

One of the technologies on your list is lab-grown meat, which is still very tentative and expensive. Why did it make the cut?
Part of the reason I picked it is to remind people that clean energy does not solve climate change. Only about a quarter of emissions come from electricity generation. This is a category that people weren't paying much attention to as a greenhouse-gas problem. And yet I think the path to solve it is clearer than in, say, cement or steel or other materials.

Another of your picks is the reinvented toilet, which you call the biggest advance in sanitation in 200 years. Why?
Building sewers, using clean water, having a processing plant—that's the paradigm in rich countries. In low-income countries, the capital cost of a sewer system is just unattainable. This toilet takes the human waste, liquid and solid, and in most cases does some type of separation. The solids you can essentially burn. The liquids you can filter. That's a huge effect on quality of life, in terms of both disgust and disease, in an increasingly urbanized world. The Gates

> “Part of the reason I picked lab-grown meat is to remind people that clean energy does not solve climate change.”

Foundation has given out $200 million in grants to try to get this technology going. It’s not there yet.

Three of your picks are about reducing greenhouse-gas emissions. You lead a $1 billion investment fund, Breakthrough Energy Ventures. But it feels like there are already a lot of technological solutions to climate change. Do we really need more? Isn’t the biggest problem political?
No, the problems are when you say to India, “Provide electricity to everyone to have things we take for granted—heating, air conditioning.” Their path is to build more coal plants. That’s the cheapest form of electricity for them. In France they were asked to pay a 5% increase on their diesel price, and even that was unacceptable.

The politics is where you decide how much you’re going to put into basic research or how you’re going to make things attractive for innovative companies. But if we freeze technology today, you will live in a 4 °C warmer world in the future, guaranteed.

One of those picks is nuclear fusion. That’s something that’s always seemed just around the corner. What makes you optimistic about it?
The company that Breakthrough put money into, Commonwealth Fusion Systems—the methods they’re using allow you to get a dramatic reduction in the size and therefore the capital cost. It’s very impressive. There are over 10 companies pursuing fusion in different ways. Most of them will not work. But these projects certainly will make a big contribution. So I think it’s important we back fusion.

China is becoming a technology superpower. How do you think that will play out as fear about its power gets entrenched?
The idea that they’re starting to be innovative—that is good for the world.

Like most middle-income countries, they’re more than willing to do big projects. Think of the US in the ’50s and ’60s, Japan in the ’70s and ’80s, Korea in the ’80s and ’90s. Your technological capability gets really strong, and you’re willing to go out and do very, very ambitious things.

For the US, it’s good to have a sense that we have to renew our edge. In the ’70s and ’80s, when we were like, “Oh jeez, has Japan figured things out we haven’t,” we renewed our commitment to basic research. In fact, Japan was never going to overtake us in terms of scientific innovation. But I do think that was healthy for us.

A

FORGET DRIVERLESS CARS AND FACE RECOGNITION —

’s big idea: Reinvent how we invent

THE BIGGEST IMPACT OF ARTIFICIAL INTELLIGENCE WILL BE REINVIGORATING HOW WE DO RESEARCH.

BY DAVID ROTMAN

Regina Barzilay’s office at MIT affords a clear view of the Novartis Institutes for Biomedical Research. Amgen’s drug discovery group is a few blocks beyond that. Until recently, Barzilay, one of the world’s leading researchers in artificial intelligence, hadn’t given much thought to these nearby buildings full of chemists and biologists. But as AI and machine learning began to perform ever more impressive feats in image recognition and language comprehension, she began to wonder: could it also transform the task of finding new drugs?

Drug discovery is a hugely expensive and often frustrating process. Medicinal chemists must guess which compounds might make good medicines, using their knowledge of how a molecule’s structure affects its properties. They synthesize and test countless variants, and most are failures. “Coming up with new molecules is still an art, because you have such a huge space of possibilities,” says Barzilay. “It takes a long time to find good drug candidates.”

The problem is that human researchers can explore only a tiny slice of what is possible. It’s estimated that there are as many as 10^{60} potentially drug-like molecules—more than the number of atoms in the solar system. But traversing seemingly unlimited possibilities is what machine learning is good at. Trained on large databases of existing molecules and their properties, the programs can explore all possible related molecules.

Machine learning is already getting as good as all but the most expert chemists at figuring out how to synthesize a compound and predicting its properties—two essential tasks in drug discovery. What Barzilay and others are now doing is creating deep-learning algorithms that can imagine entirely novel molecules with desirable properties—new “lead” compounds for chemists to tweak and test.

By speeding up this critical step, deep learning could offer far more opportunities for chemists to pursue, making drug discovery much quicker. One advantage: machine learning’s often quirky imagination. “Maybe it will go in a different direction that a human wouldn’t go in,” says Angel Guzman-Perez, a drug researcher at Amgen who is working with Barzilay. “It thinks differently.”

Others are using machine learning to try to invent new materials for clean-tech applications. Among the items on the wish list are improved batteries for storing power on the electric grid and organic solar cells, which could be far cheaper to make than today’s bulky silicon-based ones.

Such breakthroughs have become harder and more expensive to attain as chemistry, materials science, and drug discovery have grown mind-bogglingly complex and saturated with data. Even as the pharmaceutical and biotech industries pour money into research, the number of new drugs based on novel molecules has been flat over the last few decades. And we’re still stuck with lithium-ion batteries that date to the early 1990s and designs for silicon solar cells that are also decades old.

The complexity that has slowed progress in these fields is where deep learning excels. Searching through multidimensional space to come up with valuable predictions is “AI’s sweet spot,” says Ajay

Agrawal, an economist at the Rotman School of Management in Toronto and author of the best-selling *Prediction Machines: The Simple Economics of Artificial Intelligence*.

In a recent paper, economists at MIT, Harvard, and Boston University argued that AI's greatest economic impact could come from its potential as a new "method of invention" that ultimately reshapes "the nature of the innovation process and the organization of R&D." Iain Cockburn, a BU economist and coauthor of the paper, says: "New methods of invention with wide applications don't come by very often, and if our guess is right, AI could dramatically change the cost of doing R&D in many different fields." Much of innovation involves making predictions based on data. In such tasks, Cockburn adds, "machine learning could be much faster and cheaper by orders of magnitude."

In other words, AI's chief legacy might not be driverless cars or image search or even Alexa's ability to take orders, but its ability to come up with new ideas to fuel innovation itself.

The rising price of big ideas

It is taking more researchers and money to find productive new ideas, according to economists at Stanford and MIT. That's a likely factor in the overall sluggish growth in the US and Europe in recent decades. The graph below shows the pattern for the overall economy, highlighting US total factor productivity (by decade average and for 2000–2014)—a measure of the contribution of innovation—versus the number of researchers. Similar patterns hold for specific research areas.

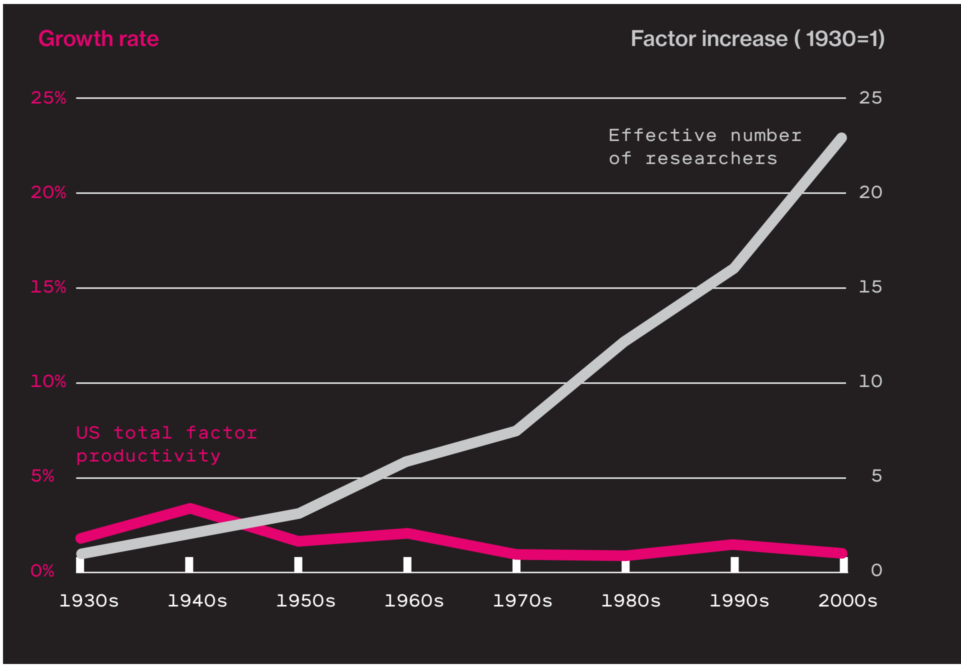

IDEAS ARE GETTING EXPENSIVE

Late last year, Paul Romer won the economics Nobel Prize for work done during the late 1980s and early 1990s that showed how investments in new ideas and innovation drive robust economic growth. Earlier economists had noted the connection between innovation and growth, but Romer provided an exquisite explanation for how it works. In the decades since, Romer's conclusions have been the intellectual inspiration for many in Silicon Valley and help account for how it has attained such wealth.

But what if our pipeline of new ideas is drying up? Economists Nicholas Bloom and Chad Jones at Stanford, Michael Webb, a graduate student at the university, and John Van Reenen at MIT looked at the problem in a recent paper called "Are ideas getting harder to find?" (Their answer was "Yes.") Looking at drug discovery, semiconductor research, medical innovation, and efforts to improve crop yields, the economists found a common story: investments in research are climbing sharply, but the payoffs are staying constant.

From an economist's perspective, that's a productivity problem: we're paying more for a similar amount of output. And the numbers look bad. Research productivity—the number of researchers it takes to produce a given result—is declining by around 6.8% annually for the task of extending Moore's Law, which requires that we find ways to pack ever more and smaller components on a semiconductor chip in order to keep making computers faster and more powerful. (It takes more than 18 times as many researchers to double chip density today as it did in the early 1970s, they found.) For improving seeds, as measured by crop yields, research productivity is dropping by around 5% each year. For the US economy as a whole, it is declining by 5.3%.

Any negative effect of this decline has been offset, so far, by the fact that we're putting more money and people into research. So we're still doubling the number of transistors on a chip every two years, but only because we're dedicating far more people to the problem. We'll have to double our investments in research and development over the next 13 years just to keep treading water.

It could be, of course, that fields like crop science and semiconductor research are getting old and the opportunities for innovation are shriveling up. However, the researchers also found that overall growth tied to innovation in the economy was slow. Any investments in new areas, and any inventions they have generated, have failed to change the overall story.

The drop in research productivity appears to be a decades-long trend. But it is particularly worrisome to economists now because we've seen an overall slowdown

SOURCE: BLOOM, JONES, VAN REENEN, AND WEBB

Corn

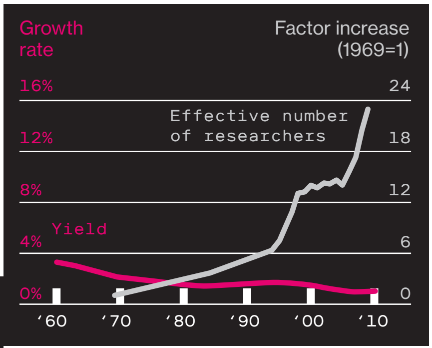

New molecular drug entities

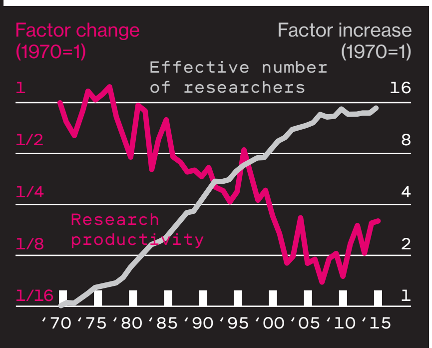

Moore's Law

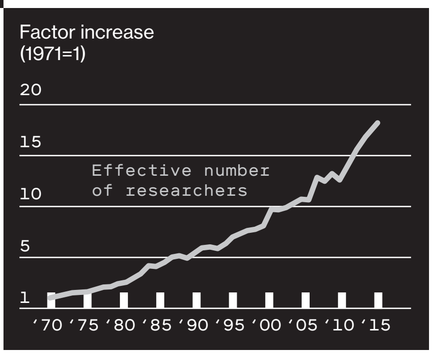

in economic growth since the mid-2000s. At a time of brilliant new technologies like smartphones, driverless cars, and Facebook, growth is sluggish, and the portion of it attributed to innovation—called total factor productivity—has been particularly weak.

The lingering effects of the 2008 financial collapse could be hampering growth, says Van Reenen, and so could continuing political uncertainties. But dismal research productivity is undoubtedly a contributor. And he says that if the decline continues, it could do serious damage to future prosperity and growth.

It makes sense that we've already picked much of what some economists like to call the "low-hanging fruit" in terms of inventions. Could it be that the only fruit left is a few shriveled apples on the farthest branches of the tree? Robert Gordon, an economist at Northwestern University, has been a strong proponent of that view. He says we're unlikely to match the flourishing of discovery that marked the late 19th and early 20th centuries, when inventions such as electric light and power and the internal-combustion engine led to a century of unprecedented prosperity.

If Gordon is right, and there are fewer big inventions left, we're doomed to a dismal economic future. But few economists think that's the case. Rather, it makes sense that big new ideas are out there; it's just getting more expensive to find them as the science becomes increasingly complex. The chances that the next penicillin will just fall into our laps are slim. We'll need more and more researchers to make sense of the advancing science in fields like chemistry and biology.

It's what Ben Jones, an economist at Northwestern, calls "the burden of knowledge." Researchers are becoming more specialized, making it necessary to form larger—and more expensive—teams to solve problems. Jones's research shows that the age at which scientists reach their peak productivity is going up: it takes them longer to gain the expertise they need. "It's an innate by-product of the exponential growth of knowledge," he says.

"A lot of people tell me our findings are depressing, but I don't see it that way," says Van Reenen. Innovation might be more difficult and expensive, but that, he says, simply points to the need for policies, including tax incentives, that will encourage investments into more research.

"As long as you put resources into R&D, you can maintain healthy productivity growth," says Van Reenen. "But we have to be prepared to spend money to do it. It doesn't come free."

GIVING UP ON SCIENCE

Can an AI creatively solve the kinds of problems that such innovation requires? Some experts are now convinced that it can, given the kinds of advances shown off by the game-playing machine AlphaGo.

AlphaGo mastered the ancient game of Go, beating the reigning champion, by studying the nearly unlimited possible moves in a game that has been played for several thousand years by humans relying heavily on intuition. In doing so, it sometimes came up with winning strategies that no human player had thought to try. Likewise, goes the thinking, deep-learning programs trained on large amounts of experimental data and chemical literature could come up with novel compounds that scientists never imagined.

Might an AlphaGo-like breakthrough help the growing armies of researchers poring over ever-expanding scientific data? Could AI make basic research faster and more productive, reviving areas that have become too expensive for businesses to pursue?

MIGHT AN ALPHAGO-LIKE BREAKTHROUGH HELP THE GROWING ARMIES OF RESEARCHERS PORING OVER EVER-EXPANDING SCIENTIFIC DATA?

The last several decades have seen a massive upheaval in our R&D efforts. Since the days when AT&T's Bell Labs and Xerox's PARC produced world-changing inventions like the transistor, solar cells, and laser printing, most large companies in the US and other rich economies have given up on basic research. Meanwhile, US federal R&D investments have been flat, particularly for fields other than life sciences. So while we continue to increase the number of researchers overall and to turn incremental advances into commercial opportunities, areas that require long-term research and a grounding in basic science have taken a hit.

The invention of new materials in particular has become a commercial backwater. That has held back needed innovations in clean tech—stuff like better batteries, more efficient solar cells, and catalysts to make fuels directly from sunlight and carbon dioxide (think artificial photosynthesis). While the prices of solar panels and batteries are falling steadily, that's largely because of improvements in manufacturing and economies of scale, rather than fundamental advances in the technologies themselves.

It takes an average of 15 to 20 years to come up with a new material, says Tonio Buonassisi, a mechanical engineer at MIT who is working with a team of scientists in Singapore to speed up the process. That's far too long for most businesses. It's impractical even for many academic groups. Who wants to spend years on a material that may or may not work? This is why venture-backed startups, which have generated much of the innovation in software and even biotech, have long given up on clean tech: venture capitalists generally need a return within seven years or sooner.

"A 10x acceleration [in the speed of materials discovery] is not only possible, it is necessary," says Buonassisi, who runs a photovoltaic research lab at MIT. His goal, and that of a loosely connected network of fellow scientists, is to use AI and machine learning to get that 15-to-20-year time frame down to around two to five years by attacking the various bottlenecks in the lab, automating as much of the process as possible. A faster process gives the scientists far more potential solutions to test, allows them to find dead ends in hours rather than months, and helps optimize the materials. "It transforms how we think as researchers," he says.

It could also make materials discovery a viable business pursuit once again. Buonassisi points to a chart showing the time it took to develop various technologies. One of the columns labeled "lithium-ion batteries" shows 20 years. Another, much shorter column is labeled "novel solar cell"; at the top is "2030 climate target." The point is clear: we can't wait another 20 years for the next breakthrough in clean-tech materials.

THE AI-DRIVEN LAB

"Come to a free land": that is how Alán Aspuru-Guzik invites a US visitor to his Toronto lab these days. In 2018 Aspuru-Guzik left his tenured position as a Harvard chemistry professor, moving with his family to Canada. His decision was driven by a strong distaste for President Donald Trump and his policies, particularly those on immigration. It didn't hurt, however, that Toronto is rapidly becoming a mecca for artificial-intelligence research.

As well as being a chemistry professor at the University of Toronto, Aspuru-Guzik also has a position at the Vector Institute for Artificial Intelligence. It's the AI center cofounded by Geoffrey Hinton, whose pioneering work on deep learning and neural networks is largely credited with jump-starting today's boom in AI.

THE IDEA IS TO INFUSE ARTIFICIAL INTELLIGENCE AND AUTOMATION INTO ALL THE STEPS OF MATERIALS RESEARCH AND DRUG DISCOVERY.

In a notable 2012 paper, Hinton and his coauthors demonstrated that a deep neural network, trained on a huge number of pictures, could identify a mushroom, a leopard, and a dalmatian dog. It was a remarkable breakthrough at the time, and it quickly ushered in an AI revolution using deep-learning algorithms to make sense of large data sets. Researchers rapidly found ways to use such neural networks to help driverless cars navigate and to spot faces in a crowd. Others modified the deep-learning tools so that they could train themselves; among these tools are GANs (generative adversarial networks), which can fabricate images of scenes and people that never existed.

In a 2015 follow-up paper, Hinton provided clues that deep learning could be used in chemistry and materials research. His paper touted the ability of neural network to discover "intricate structures in high-dimensional data"—in other words, the same networks that can navigate through millions of images to find, say, a dog with spots could sort through millions of molecules to identify one with certain desirable properties.

Energetic and bubbling with ideas, Aspuru-Guzik is not the type of scientist to patiently spend two decades figuring out whether a material will work. And he has quickly adapted deep learning and neural networks to attempt to reinvent materials discovery. The idea is to infuse artificial intelligence and automation into all the steps of materials research: the initial design and synthesis of a material, its testing and analysis, and finally the multiple refinements that optimize its performance.

On a freezing cold day early this January, Aspuru-Guzik has his hat pulled tightly down over his ears but otherwise seems oblivious to the bitter Canadian weather. He has other things on his mind. For one thing, he's still waiting for the delivery of a $1.2 million robot, now on a ship from Switzerland, that will be the centerpiece for the automated, AI-driven lab he has envisioned.

In the lab, deep-learning tools like GANs and their cousin, a technique called autoencoder, will imagine promising new materials and figure out how to make them. The robot will then make the compounds; Aspuru-Guzik wants to create an affordable automated system that would be able to spit out new molecules on demand. Once the materials have been made, they can be analyzed with instruments such as a mass spectrometer. Additional machine-learning tools will make sense of that data and "diagnose" the material's properties. These insights

AI startups in drugs and materials

	1 Atomwise	2 Kebotix	3 Deep Genomics
What they do	Use neural networks to search through large databases to find small drug-like molecules that bind to targeted proteins.	Develop a combination of robotics and AI to speed up the discovery and development of new materials and chemicals.	Use artificial intelligence to search for oligonucleotide molecules to treat genetic diseases.
Why it matters	Identifying such molecules with desirable properties, such as potency, is a critical first step in drug discovery.	It takes more than a decade to develop a material. Cutting that time could help us tackle problems such as climate change.	Oligonucleotide treatments hold promise against a range of diseases, including neurodegenerative and metabolic disorders.

will then be used to further optimize the materials, tweaking their structures. And then, Aspuru-Guzik says, "AI will select the next experiment to make, closing the loop."

Once the robot is in place, Aspuru-Guzik expects to make some 48 novel materials every two days, drawing on the machine-learning insights to keep improving their structures. That's one promising new material every hour, an unprecedented pace that could completely transform the lab's productivity.

It's not all about simply dreaming up "a magical material," he says. To really change materials research, you need to attack the entire process: "What are the bottlenecks? You want AI in every piece of the lab." Once you have a proposed structure, for example, you still need to figure out how to make it. It can take weeks to months to solve what chemists call "retrosynthesis"—working backwards from a molecular structure to figure out the steps needed to synthesize such a compound. Another bottleneck comes in making sense of the reams of data produced by analytic equipment. Machine learning could speed up each of those steps.

What motivates Aspuru-Guzik is the threat of climate change, the need for improvements in clean tech, and the essential role of materials in producing such advances. His own research is looking at novel organic electrolytes for flow batteries, which can be used to store excess electricity from power grids and pump it back in when it's needed, and at organic solar cells that would be far cheaper than silicon-based ones. But if his design for a self-contained, automated chemical lab works, he suggests, it could make chemistry far more accessible to almost anyone. He calls it the "democratization of materials discovery."

A few blocks away, Ajay Agrawal runs the Creative Destruction Lab at the Rotman business school. The program has spawned more than 200 startups since its inception in 2012. Many have originated with computer science students who wander in, looking to apply machine learning to everything from spotting credit card fraud to identifying tumors in medical images. These days, though, Agrawal is intensely focused on how these same AI tools could be applied to accelerating scientific research.

"This is where the action is," he says. "AIs that drive cars, AIs that improve medical diagnostics, AIs for personal shopping—the economic growth from AIs applied to scientific research may swamp the economic impact from all those other AIs combined."

The Vector Institute, Toronto's magnet for AI research, sits less than a mile away. From the windows of the large open office space, you can look across at Ontario's parliament building. The proximity of experts in AI, chemistry, and business to the province's seat of government in downtown Toronto isn't a coincidence. There's a strong belief among many in the city that AI will transform business and the economy, and increasingly, some are convinced it will radically change how we do science.

WILL SCIENTISTS BUY IN?

Still, if it is do that, a first step is convincing scientists it is worthwhile.

Amgen's Guzman-Perez says many of his peers in medicinal chemistry are skeptical. Over the last few decades the field has seen a series of supposedly revolutionary technologies, from computational design to combinatorial chemistry and high-throughput screening, that have automated the rapid production and testing of multiple molecules. Each has proved somewhat helpful but limited. None, he says, "magically get you a new drug."

It's too early to know for sure whether deep learning could finally be the game-changer, he acknowledges, "and it's hard to know the time frame." But he takes encouragement from the speed at which AI has transformed image recognition and other search tasks. "Hopefully, it could happen in chemistry," he says.

We're still waiting for the AlphaGo moment in chemistry and materials—for deep-learning algorithms to outwit the most accomplished human in coming up with a new drug or material. But just as AlphaGo won with a combination of uncanny strategy and an inhuman imagination, today's latest AI programs could soon prove themselves in the lab.

And that has some scientists dreaming big. The idea, says Aspuru-Guzik, is to use AI and automation to reinvent the lab with tools such as the $30,000 molecular printer he hopes to build. It will then be up to scientists' imagination—and that of AI—to explore the possibilities.

David Rotman is editor at large at MIT Technology Review.


MIT Technology Review
EmTech Next
June 11–12, 2019
MIT Media Lab
Cambridge, MA
When robots are your colleagues, which human skills will still matter?
EmTech Next is a two-day exploration of how advances in AI and other digital technologies are transforming the future of work and the economy.
Register at
EmTechNext.com/2019

Meet the pioneers shaping the future of work, including the connected factory, human/robot collaboration, and more.

Henny Admoni

Carnegie Mellon University

David Autor

MIT

Jit Kee Chin

Suffolk

Moustapha Cisse

Google

Shelley Peterson

Lockheed Martin

Can machines be truly creative?

AlphaZero, a computer program that taught itself to be a chess grandmaster in a few hours, exhibits "the essence of creativity," says its creator.

By Will Knight
Portrait by Geordie Wood

David Silver invented something that might be more inventive than he is.

Silver was the lead researcher on AlphaGo, a computer program that learned to play Go—a famously tricky game that exploits human intuition rather than clear rules of play—by studying games played by humans.

Silver's latest creation, AlphaZero, learns to play board games including Go, chess, and Shogu by practicing against itself. Through millions of practice games, AlphaZero discovers strategies that it took humans millennia to develop.

So could AI one day solve problems that human minds never could? I spoke to Silver at his London office at DeepMind, now owned by Alphabet.

In one famous game against possibly the best Go player ever, AlphaGo made a brilliant move that human observers initially thought was a mistake. Was it being creative in that moment?
"Move 37," as it became known, surprised everyone, including the Go community and us, its makers. It was something outside of the expected way of playing Go that humans had figured out over thousands of years. To me this is an example of something being creative.

Since AlphaZero doesn't learn from humans, is it even more creative?
When you have something learning by itself, that's building up its own knowledge completely from scratch, it's almost the essence of creativity.

AlphaZero has to figure out everything for itself. Every single step is a creative leap. Those insights are creative because they weren't given to it by humans. And those leaps continue until it is something that is beyond our abilities and has the potential to amaze us.

You've had AlphaZero play against the top conventional chess engine, Stockfish. What have you learned?
Stockfish has this very sophisticated search engine, but at the heart of it is this module that says, "According to humans, this is a good position or a bad position." So humans are really deeply in the loop there. It's hard for it to break away and understand a position that's fundamentally different.

AlphaZero learns to understand positions for itself. There was one beautiful game we were just looking at where it actually gives up four pawns in a row, and it even tries to give up a fifth pawn. Stockfish thinks it's winning fantastically, but AlphaZero is really happy. It's found a way to understand the position which is unthinkable according to the norms of chess. It understands it's better to have the position than the four pawns.

Does AlphaZero suggest AI will play a role in future scientific innovation?
Machine learning has been dominated by an approach called supervised learning, which means you start off with everything that humans know, and you try to distill that into a computer program that does things in just the same way. The beauty of this new approach, reinforcement learning, is that the system learns for itself, from first principles, how to achieve the goals we set it. It's like a million mini-discoveries, one after another, that build up this creative way of thinking. And if you can do that, you can end up with something that has immense power, immense ability to solve problems, and which can hopefully lead to big breakthroughs.

Are there aspects of human creativity that couldn't be automated?
If we think about the capabilities of the human mind, we're still a long way away from achieving that. We can achieve results in specialized domains like chess and Go with a massive amount of computer power dedicated to that one task. But the human mind is able to radically generalize to something different. You can change the rules of the game, and a human doesn't need another 2,000 years to figure out how she should play.

I would say that maybe the frontier of AI at the moment—and where we'd like to go—is to increase the range and the flexibility of our algorithms to cover the full gamut of what the human mind can do. But that's still a long way off.

How might we get there?
I'd like to preserve this idea that the system is free to create without being constrained by human knowledge.

A baby doesn't worry about its career, or how many kids it's going to have. It is playing with toys and learning manipulation skills. There's an awful lot to learn about the world in the absence of a final goal. The same can and should be true of our systems.


호프
모
텔

What computers can't create

Why creativity is, and always will be, a human endeavor

by Sean Dorrance Kelly

On March 31, 1913, in the Great Hall of the Musikverein concert house in Vienna, a riot broke out in the middle of a performance of an orchestral song by Alban Berg. Chaos descended. Furniture was broken. Police arrested the concert's organizer for punching Oscar Straus, a little-remembered composer of operettas. Later, at the trial, Straus quipped about the audience's frustration. The punch, he insisted, was the most harmonious sound of the entire evening. History has rendered a different verdict: the concert's conductor, Arnold Schoenberg, has gone down as perhaps the most creative and influential composer of the 20th century.

You may not enjoy Schoenberg's dissonant music, which rejects conventional tonality to arrange the 12 notes of the scale according to rules that don't let any predominate. But he changed what humans understand music to be. This is what makes him a genuinely creative and innovative artist. Schoenberg's techniques are now integrated seamlessly into everything from film scores and Broadway musicals to the jazz solos of Miles Davis and Ornette Coleman.

Creativity is among the most mysterious and impressive achievements of human existence. But what is it?

rs

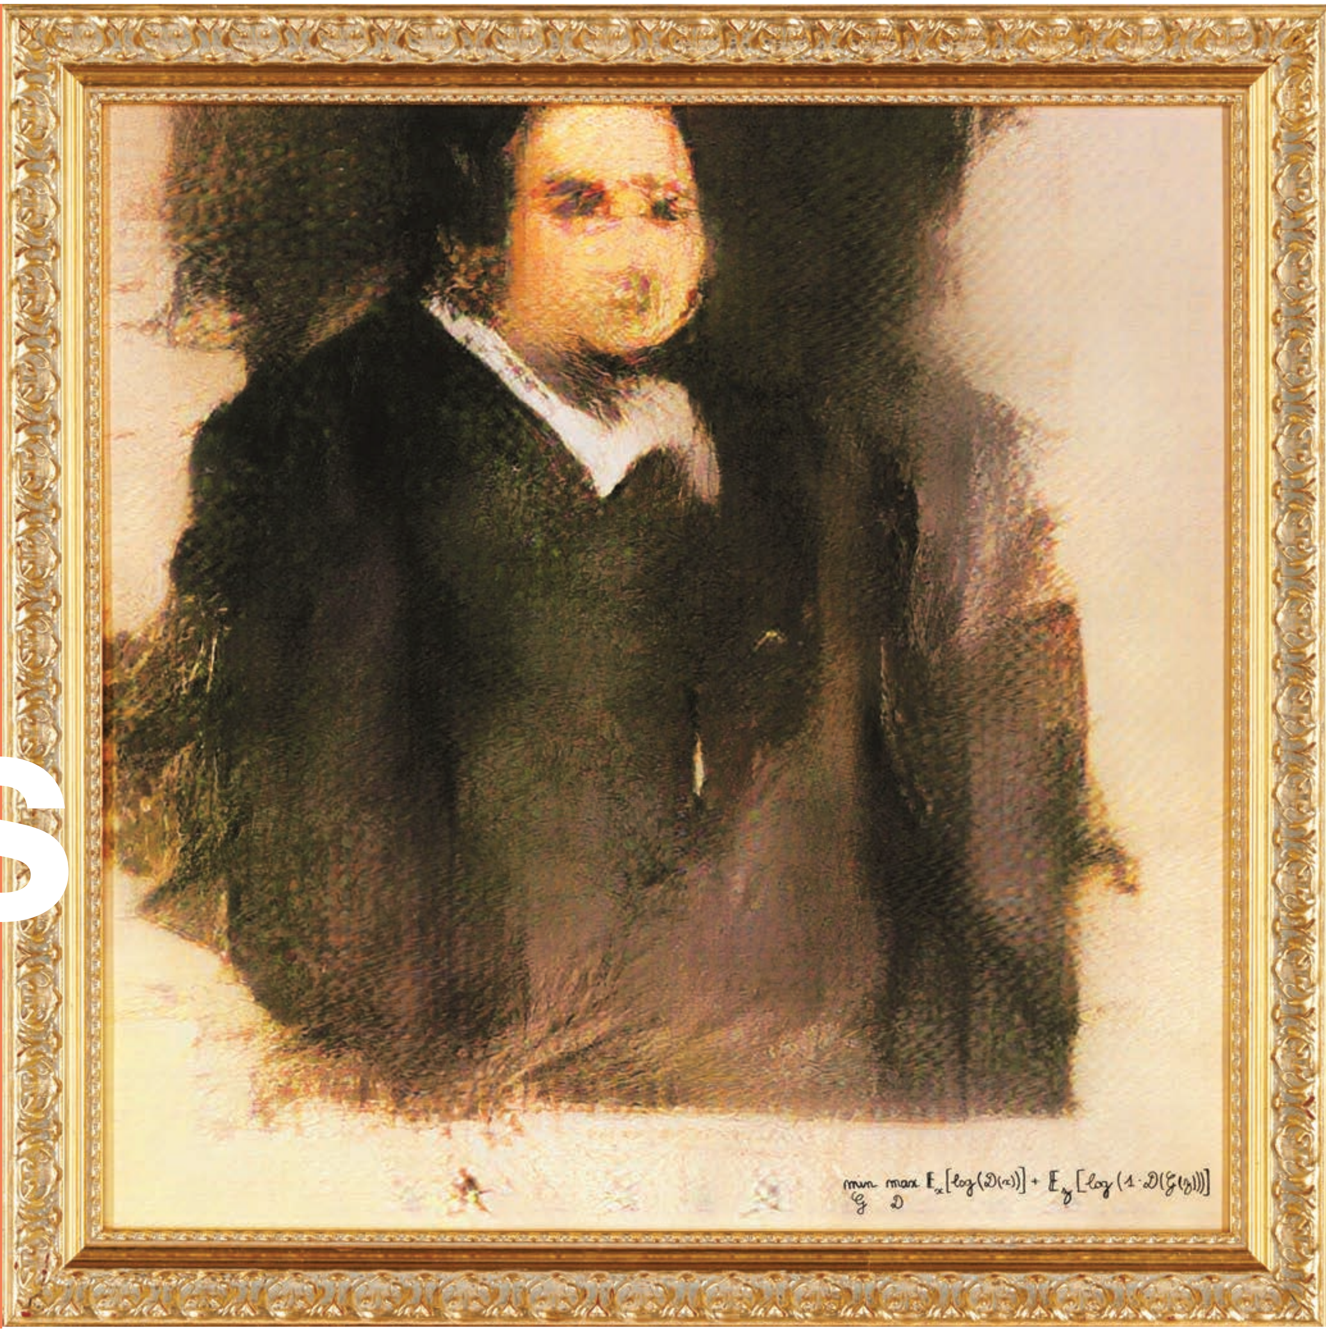

Portrait of Edmond Belamy (2018), created with AI algorithms called GANs by Parisian art collective Obvious, sold for $432,500.

COURTESY OF THE ARTISTS

Creativity is not just novelty. A toddler at the piano may hit a novel sequence of notes, but they're not, in any meaningful sense, creative. Also, creativity is bounded by history: what counts as creative inspiration in one period or place might be disregarded as ridiculous, stupid, or crazy in another. A community has to accept ideas as *good* for them to count as creative.

As in Schoenberg's case, or that of any number of other modern artists, that acceptance need not be universal. It might, indeed, not come for years—sometimes creativity is mistakenly dismissed for generations. But unless an innovation is eventually accepted by some community of practice, it makes little sense to speak of it as creative.

Advances in artificial intelligence have led many to speculate that human beings will soon be replaced by machines in every domain, including that of creativity. Ray Kurzweil, a futurist, predicts that by 2029 we will have produced an AI that can pass for an average educated human being. Nick Bostrom, an Oxford philosopher, is more circumspect. He does not give a date but suggests that philosophers and mathematicians defer work on fundamental questions to "superintelligent" successors, which he defines as having "intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest."

Human creative achievement, because of the way it is socially embedded, will not succumb to advances in artificial intelligence.

Both believe that once human-level intelligence is produced in machines, there will be a burst of progress—what Kurzweil calls the "singularity" and Bostrom an "intelligence explosion"—in which machines will very quickly supersede us by massive measures in every domain. This will occur, they argue, because superhuman achievement is the same as ordinary human achievement except that all the relevant computations are performed much more quickly, in what Bostrom dubs "speed superintelligence."

So what about the highest level of human achievement—creative innovation? Are our most creative artists and thinkers about to be massively surpassed by machines?

No.

Human creative achievement, because of the way it is socially embedded, will not succumb to advances in artificial intelligence. To say otherwise is to misunderstand both what human beings are and what our creativity amounts to.

This claim is not absolute: it depends on the norms that we allow to govern our culture and our expectations of technology. Human beings have, in the past, attributed great power and genius even to lifeless totems. It is entirely possible that we will come to *treat* artificially intelligent machines as so vastly superior to us that we will naturally attribute creativity to them. Should that happen, it will not be because machines have outstripped us. It will be because we will have denigrated ourselves.

Also, I am primarily talking about machine advances of the sort seen recently with the current deep-learning paradigm, as well as its computational successors. Other paradigms have governed AI research in the past. These have already failed to realize their promise. Still other paradigms may come in the future, but if we speculate that some notional future AI whose features we cannot meaningfully describe will accomplish wondrous things, that is mythmaking, not reasoned argument about the possibilities of technology.

Creative achievement operates differently in different domains. I cannot offer a complete taxonomy of the different kinds of creativity here, so to make the point I will sketch an argument involving three quite different examples: music, games, and mathematics.

Music to my ears

Can we imagine a machine of such superhuman creative ability that it brings about changes in what we understand music to be, as Schoenberg did?

That's what I claim a machine cannot do. Let's see why.

Computer music composition systems have existed for quite some time. In 1965, at the age of 17, Kurzweil himself, using a precursor of the pattern recognition systems that characterize deep-learning algorithms today, programmed a computer to compose recognizable music. Variants of this technique are used today. Deep-learning algorithms have been able to take as input a bunch of Bach chorales, for instance, and compose music so characteristic of Bach's style that it fools even experts into thinking it is original. This is mimicry. It is what an artist does as an apprentice:

In Imaginary Landscape (2018), Nao Tokui uses a machine-learning algorithm to create panoramas from images found in Google Street View and complements them with soundscapes created with artificial neural networks.

COURTESY OF THE ARTIST

copy and perfect the style of others instead of working in an authentic, original voice. It is not the kind of musical creativity that we associate with Bach, never mind with Schoenberg's radical innovation.

So what do we say? Could there be a machine that, like Schoenberg, invents a whole new way of making music? Of course we can imagine, and even make, such a machine. Given an algorithm that modifies its own compositional rules, we could easily produce a machine that makes music as different from what we now consider good music as Schoenberg did then.

But this is where it gets complicated.

We count Schoenberg as a creative innovator not just because he managed to create a new way of composing music but because people could see in it a vision of what the world should be. Schoenberg's vision involved the spare, clean, efficient minimalism of modernity. His innovation was not just to find a new algorithm for composing music; it was to find a way of thinking about what music *is* that allows it to speak to *what is needed now*.

Some might argue that I have raised the bar too high. Am I arguing, they will ask, that a machine needs some mystic, unmeasurable sense of what is socially necessary in order to count as creative? I am not—for two reasons.

First, remember that in proposing a new, mathematical technique for musical composition, Schoenberg changed our understanding of what music is. It is only creativity of this tradition-defying sort that requires some kind of social sensitivity. Had listeners not experienced his technique as capturing the anti-traditionalism at the heart of the radical modernity emerging in early-20th-century Vienna, they might not have heard it as something of aesthetic worth. The point here is that radical creativity is not an "accelerated" version of quotidian creativity. Schoenberg's achievement is not a faster or better version of the type of creativity demonstrated by Oscar Straus or some other average composer: it's fundamentally different in kind.

Second, my argument is not that the creator's responsiveness to social necessity must be conscious for the work to meet the standards of genius. I am arguing instead that *we must be able to interpret the work as responding that way.* It would be a mistake to interpret a machine's composition as part of such a vision of the world. The argument for this is simple.

Claims like Kurzweil's that machines can reach human-level intelligence assume that to have a human mind is just to have a human brain that follows some set of computational algorithms—a view called computationalism. But though algorithms can have moral implications, they are not themselves moral agents. We can't count the monkey at a typewriter who accidentally types out *Othello* as a great creative playwright. If there is greatness in the product, it is only an accident. *We* may be able to see a machine's *product* as great, but if we know that the output is merely the result of some arbitrary act or algorithmic formalism, we cannot accept it as the expression of a vision for human good.

For this reason, it seems to me, nothing but another human being can properly be understood as a genuinely creative artist. Perhaps AI will someday proceed beyond its computationalist formalism, but that would require a leap that is unimaginable at the moment. We wouldn't just be looking for new algorithms or procedures that simulate human activity; we would be looking for new materials that are the basis of being human.

A molecule-for-molecule duplicate of a human being would be human in the relevant way. But we already have a way of producing such a being: it takes about nine months. At the moment, a machine can only do something much less interesting than what a person can do. It can create music in the style of Bach, for instance—perhaps even music that some experts think is better than Bach's own. But that is only because its music can be judged against a preexisting standard. What a machine cannot do is bring about changes in our standards for judging the quality of music or of understanding what music is or is not.

This is not to deny that creative artists use whatever tools they have at their disposal, and that those tools shape the sort of art they make. The trumpet helped Davis and Coleman realize their creativity. But the trumpet is not, itself, creative. Artificial-intelligence algorithms are more like musical instruments than they are like people. Taryn Southern, a former *American Idol* contestant, recently released an album where the percussion, melodies, and chords were algorithmically generated, though she wrote the lyrics and repeatedly tweaked the instrumentation algorithm until it delivered the results she wanted. In the early 1990s, David Bowie did it the other way around: he wrote the music and used a Mac app called Verbalizer to pseudorandomly recombine sentences into lyrics. Just like previous tools of the music industry—from recording devices to synthesizers to samplers and loopers—new AI tools work by stimulating and channeling the creative abilities of the human artist (and reflect the limitations of those abilities).

Anna Ridler's The Fall of the House of Usher (2017) is a 12-minute animation based on Watson and Webber's 1928 silent film. Ridler created the stills using three separate neural nets: one trained on her drawings, a second trained on drawings made of the results of the first net, and a third trained on drawings made of the results of the second.

COURTESY OF THE ARTIST

Games without frontiers

Much has been written about the achievements of deep-learning systems that are now the best Go players in the world. AlphaGo and its variants have strong claims to having created a whole new way of playing the game. They have taught human experts that opening moves long thought to be ill-conceived can lead to victory. The program plays in a style that experts describe as strange and alien. "They're how I imagine games from far in the future," Shi Yue, a top Go player, said of AlphaGo's play. The algorithm seems to be genuinely creative.

In some important sense it is. Game-playing, though, is different from composing music or writing a novel: in games there is an objective measure of success. We know we have something to learn from AlphaGo because we see it win. But that is also what makes Go a "toy domain," a simplified case that says only limited things about the world.

The most fundamental sort of human creativity changes our understanding of ourselves because it changes our understanding of what we count as good. For the game of Go, by contrast, the nature of goodness is simply not up for grabs: a Go strategy is good if and only if it wins. Human life does not generally have this feature: there is no objective measure of success in the highest realms of achievement. Certainly not in art, literature, music, philosophy, or politics. Nor, for that matter, in the development of new technologies.

In various toy domains, machines may be able to teach us about a certain very constrained form of creativity. But the domain's rules are pre-formed; the system can succeed only because it learns to play well within these constraints. Human culture and human existence are much more interesting than this. There are norms for how human beings act, of course. But creativity in the genuine sense is the ability to change those norms in some important human domain. Success in toy domains is no indication that creativity of this more fundamental sort is achievable.

It's a knockout

A skeptic might contend that the argument works only because I'm contrasting games with *artistic* genius. There are other paradigms of creativity in the scientific and mathematical realm. In these realms, the question isn't about a *vision* of the world. It is about the way things actually are.

Might a machine come up with mathematical proofs so far beyond us that we simply have to defer to its creative genius?

No.

Computers have already assisted with notable mathematical achievements. But their contributions haven't been particularly creative. Take the first major theorem proved using a computer: the four-color theorem, which states that any flat map can be colored with at most four colors in such a way that no two adjacent "countries" end up with the same one (it also applies to countries on the surface of a globe). Nearly a half-century ago, in 1976, Kenneth Appel and Wolfgang Haken at the University of Illinois published a computer-assisted proof of this theorem. The computer performed billions of calculations, checking thousands of different types of maps—so many that it was (and remains) logistically unfeasible for humans to verify that each possibility accorded with the computer's view. Since then, computers have assisted in a wide range of new proofs.

But the supercomputer is not doing anything creative by checking a huge number of cases. Instead, it is doing something boring a huge number of times. This seems like almost the opposite of creativity. Furthermore, it is so far from the kind of *understanding* we normally think a mathematical proof should offer that some experts don't consider these computer-assisted strategies mathematical proofs at all. As Thomas Tymoczko, a philosopher of mathematics, has argued, if we can't even verify whether the proof is correct, then all we are really doing is trusting in a potentially error-prone computational process.

Even supposing we do trust the results, however, computer-assisted proofs are something like the analogue of computer-assisted composition. If they give us a worthwhile product, it is mostly because of the contribution of the human being. But some experts have argued that artificial intelligence will be able to achieve more than this. Let us suppose, then, that we have the ultimate: a self-reliant machine that proves new theorems all on its own.

Could a machine like this massively surpass us in mathematical creativity, as Kurzweil and Bostrom argue? Suppose, for instance, that an AI comes up with a resolution to some extremely important and difficult open problem in mathematics.

There are two possibilities. The first is that the proof is extremely clever, and when experts in the field

Tom White uses "perception engines," algorithms that distill the data collected from thousands of photographs of common objects, to synthesize abstract shapes. He then tests and refines the results until they are recognizable by the system, as seen in Electric Fan (2018), above.

ELECTRIC FAN, COURTESY OF TOM WHITE, MAS '98, DRIB.NET

Mario Klingemann used two GANs, one trained on a data set of human poses and one trained on pornography, to render thousands of composite images. After evaluating each for pose and detail, he chose one to refine into the finished work, The Butcher's Son (2018).

COURTESY OF THE ARTIST

go over it they discover that it is correct. In this case, the AI that discovered the proof would be applauded. The machine itself might even be considered to be a creative mathematician. But such a machine would not be evidence of the singularity; it would not so outstrip us in creativity that we couldn't even understand what it was doing. Even if it had this kind of human-level creativity, it wouldn't lead inevitably to the realm of the superhuman.

Some mathematicians are like musical virtuosos: they are distinguished by their fluency in an existing idiom. But geniuses like Srinivasa Ramanujan, Emmy Noether, and Alexander Grothendieck arguably reshaped mathematics just as Schoenberg reshaped music. Their achievements were not simply proofs of long-standing hypotheses but new and unexpected forms of reasoning, which took hold not only on the strength of their logic but also on their ability to convince other mathematicians of the significance of their innovations. A notional AI that comes up with a clever proof to a problem that has long befuddled human mathematicians is akin to AlphaGo and its variants: impressive, but nothing like Schoenberg.

That brings us to the other option. Suppose the best and brightest deep-learning algorithm is set loose and after some time says, "I've found a proof of a fundamentally new theorem, but it's too complicated for even your best mathematicians to understand."

This isn't actually possible. A proof that not even the best mathematicians can understand doesn't really count as a proof. *Proving* something implies that you are proving it *to someone*. Just as a musician has to persuade her audience to accept her aesthetic concept of what is good music, a mathematician has to persuade other mathematicians that there are good reasons to believe her vision of the truth. To count as a valid proof in mathematics, a claim must be understandable and endorsable by some independent set of experts who are in a good position to understand it. If the experts who should be able to understand the proof can't, then the community refuses to endorse it as a proof.

For this reason, mathematics is more like music than one might have thought. A machine could not surpass us massively in creativity because either its achievement would be understandable, in which case it would not *massively* surpass us, or it would not be understandable, in which case we could not count it as making any creative advance at all.

The eye of the beholder

Engineering and applied science are, in a way, somewhere between these examples. There is something like an objective, external measure of success. You can't "win" at bridge building or medicine the way you can at chess, but one can see whether the bridge falls down or the virus is eliminated. These objective criteria come into play only once the domain is fairly well specified: coming up with strong, lightweight materials, say, or drugs that combat particular diseases. An AI might help in drug discovery by, in effect, doing the same thing as the AI that composed what sounded like a well-executed Bach cantata or came up with a brilliant Go strategy. Like a microscope, telescope, or calculator, such an AI is properly understood as a tool that enables human discovery—not as an autonomous creative agent.

The capacity for genuine creativity, the kind of creativity that updates our understanding of the nature of being, is at the ground of what it is to be human.

It's worth thinking about the theory of special relativity here. Albert Einstein is remembered as the "discoverer" of relativity—but not because he was the first to come up with equations that better describe the structure of space and time. George Fitzgerald, Hendrik Lorentz, and Henri Poincaré, among others, had written down those equations before Einstein. He is acclaimed as the theory's discoverer because he had an original, remarkable, and true understanding of what the equations *meant* and could convey that understanding to others.

For a machine to do physics that is in any sense comparable to Einstein's in creativity, it must be able to persuade other physicists of the worth of its ideas at least as well as he did. Which is to say, we would have to be able to accept its proposals as aiming to *communicate their own validity to us*. Should such a machine ever come into being, as in the parable of Pinocchio, we would have to treat it as we would a human being. That means, among other things, we would have to attribute to it not only intelligence but whatever dignity and moral worth is appropriate to human beings as well. We are a long way off from

this scenario, it seems to me, and there is no reason to think the current computationalist paradigm of artificial intelligence—in its deep-learning form or any other—will ever move us closer to it.

Creativity is one of the defining features of human beings. The capacity for genuine creativity, the kind of creativity that updates our understanding of the nature of being, that changes the way we understand what it is to be beautiful or good or true—that capacity is at the ground of what it is to be human. But this kind of creativity depends upon our valuing it, and caring for it, as such. As the writer Brian Christian has pointed out, human beings are starting to act less like beings who value creativity as one of our highest possibilities, and more like machines themselves.

How many people today have jobs that require them to follow a predetermined script for their conversations? How little of what we know as real, authentic, creative, and open-ended human conversation is left in this eviscerated charade? How much is it like, instead, the kind of rule-following that a machine can do? And how many of us, insofar as we allow ourselves to be drawn into these kinds of scripted performances, are eviscerated as well? How much of our day do we allow to be filled with effectively machine-like activities—filling out computerized forms and questionnaires, responding to click-bait that works on our basest, most animal-like impulses, playing games that are designed to optimize our addictive response?

We are in danger of this confusion in some of the deepest domains of human achievement as well. If we allow ourselves to say that machine proofs we cannot understand are genuine "proofs," for example, ceding social authority to machines, we will be treating the achievements of mathematics as if they required no human understanding at all. We will be taking one of our highest forms of creativity and intelligence and reducing it to a single bit of information: yes or no.

Even if we had that information, it would be of little value to us without some understanding of the reasons underlying it. We must not lose sight of the essential character of reasoning, which is at the foundation of what mathematics is. So too with art and music and philosophy and literature. If we allow ourselves to slip in this way, to treat machine "creativity" as a substitute for our own, then machines will indeed come to seem incomprehensibly superior to us. But that is because we will have lost track of the fundamental role that creativity plays in being human.

Sean Dorrance Kelly is a philosophy professor at Harvard and coauthor of the New York Times best-selling book All Things Shining.

Puncturing dreams of drones

Unmanned aerial vehicles have been touted as a "leapfrog" solution to Africa's poor infrastructure. A researcher who studies them offers a dose of realism.

By Konstantin Kakaes
Portrait by Kate Warren

New technologies are never introduced into a vacuum. They emerge into a social, economic, and political setting and influence it in their turn. Katherine Chandler, a professor in the culture and politics program at Georgetown University, is researching drones in Africa as a study of how technology and society change together. We recently spoke with Chandler about her project.

How are drones used in Africa today?
There are a number of small-scale drone projects throughout the continent, ranging from counting wildlife to delivering vaccines to mapping islands to using drones as disaster-response technologies. One of the projects that I'm interested in is an initiative by the State University of Zanzibar. The team uses small commercial drones that can only fly for 30 or 40 minutes. So mapping Zanzibar has taken over two years.

The intention was for students to make a map that could be used for planning and natural resource management, so you would have a baseline idea of what the islands looked like if there were a hurricane, oil spill, or some other disaster. The project was not originally about resolving long-standing land claims. But part of the challenge of mapping in Zanzibar and making the information public has been figuring out how the map impacts disputes over land.

How can data gathered by drones resolve land disputes?
It's unclear how it would, or if it will. There are clearly political concerns about what this map will mean and how it's going to be used. There is a lot of information that becomes available through this high-resolution map. You can see trash dumping sites; you can see wastewater runoff; you can see where illegal building is happening. And that information changes the terms of debate.

The African Union and various international aid agencies have described drones as "transformative" for African development in general. Are they?
It's useful to think about how small an island Zanzibar is, and how long it took to carry out this particular project. When you're working in much larger spaces it becomes harder to actually cover the territory.

Take another example. Between 2016 and 2017 there was an experiment to try to integrate unmanned aircraft systems into anti-poaching efforts at Kruger National Park in South Africa. The manager in charge said that they weren't able to see any poachers by using drones and that, despite the hype around drones as an innovative new technology, drones were not capable of doing the work that was necessary to track and follow poachers, and so the project was canceled. Drones couldn't cover enough ground to gather useful information, nor were park authorities able to put the information drones gathered to good use.

There were experiments in another, much smaller, park that suggested that drones might be slightly more useful. I point this out because one of the things that I'm trying to argue is this question of scale is important when thinking about what drones can accomplish.

Fuel and battery life are a problem. Most drones right now are able to fly for no more than an hour at most. The other big limitation is payload. The amount of weight that a drone can carry is limited. This means deliveries have focused on things like blood and vaccines.

Is drone delivery a way to "leapfrog" past the need to build a better road network in much of rural Africa, where muddy roads are often impassable during rainy season?
One project that gets a lot of publicity is a venture in Rwanda by a company called Zipline to deliver blood by drone. Rwanda has been a site for huge investments by all kinds of international development organizations, and the Rwandan government is broadly interested in using drone aircraft for lots of different research projects. This has led to a vision of the country as a kind of technology hub.

But Rwanda continues to be a hugely agrarian society. How do drones fit with the day-to-day realities of most of the people living there? It is a challenge to understand who these technological investments are working for. Drones are imagined as a replacement for other forms of infrastructure, but maybe those other forms of infrastructure are actually really necessary.

It illustrates the fallacy of talking about drones as a leapfrogging technology. Thinking about how we are going to organize technologies in ways that are effectively going to serve people and communities—that's the sort of visioning that I want to see people doing.

OUR BO

OU

CELLS

DIES,

Women's health is often viewed through the lens of fertility, a bias that stymies innovation in other areas. NextGen Jane is among a vanguard of startups aiming to fix that.

By Dayna Evans

R

On an unremarkable side street in Oakland, California, a few blocks down from an animal dermatologist and just past an organic grocery store, Ridhi Tariyal and Stephen Gire are trying to change how women monitor their health. When I visited their small office in January, a garland of tampons dip-dyed in rainbow colors was strung above a computer monitor—a tongue-in-cheek reference to their work.

The tampon is a sort of totem for NextGen Jane, a startup that Tariyal and Gire founded in 2014. Their plan is to use blood squeezed from used tampons as a diagnostic tool. In that menstrual blood, they hope to find early markers of endometriosis and, ultimately, a variety of other disorders. The simplicity and ease of this method, should it work, will represent a big improvement over the present-day standard of care.

Surgeons diagnose endometriosis—an abnormal growth of endometrial tissue outside the uterus—by inserting a small camera into the pelvic cavity to look for endometrial cells in places other than the lining of the uterus, the only place they should normally grow. If wayward cells are found, the diseased tissue can often be removed on sight. But the average woman diagnosed with endometriosis has already had the disease for over a decade, which can mean years of excruciating pain.

The physical and emotional impact on women's lives is enormous. But women often believe such pain is normal, so they don't seek treatment. Delayed diagnoses by doctors relying on subjective reports of pain are also common. "I was told by my doctors that I had a 'low threshold for pain' and that I should just get used to it because there was nothing that could be done," Padma Lakshmi, a television host who founded the Endometriosis Foundation of America, said at a conference in April 2018.

A majority of endometriosis cases are never diagnosed: the most obvious symptoms can have multiple causes, and the severity of the symptoms does not correlate strongly with the severity of the underlying disorder. By some estimates, endometriosis affects 10% of reproductive-age women—roughly 200 million people.

Nevertheless, NextGen Jane did not set out to diagnose endometriosis. The company's initial focus was on fertility—because, Tariyal says, that's what venture capitalists were most interested in funding. NextGen Jane is one of hundreds of so-called femtech startups that are developing technologies intended specifically to improve women's health. Frost & Sullivan, a market research firm, predicts that femtech will be a $50 billion industry by 2025. "Women's health care," according to Frost & Sullivan, "remains largely confined to reproductive matters." According to Tariyal, this has been a major obstacle. "We wish we could go out there and say we just want to diagnose women's diseases," she told me. But investors would ask her: "Where's the money in that?"

NextGen Jane's story is a case study in how a woman's health is typically viewed through the lens of her ability to bear children—and how that ingrained bias slows innovation in medicine.

ALIENATED AND FRUSTRATED

Tariyal, who has a bachelor's degree in industrial engineering from Georgia Tech, went to work at Bank

CHRISTIE HEMM KLOK

Gire and Tariyal in their Oakland office.

of America Securities after graduation, but she hated investment banking. If she was going to grind tirelessly, she reasoned, she wanted to do something more meaningful. So she took a job as a research manager and analyst at Bristol-Myers Squibb, a pharmaceutical company. This taught her that she didn't like big companies but did love medicine. She went back to school, first getting an MBA from Harvard and then a master's in biomedical enterprise from MIT, with the goal of starting a company of her own.

As a thesis project at MIT, Tariyal tried to launch Ujala, a company that planned to test the blood of would-be partners in arranged marriages for genetic defects their offspring might inherit. It never took off. Consumer genetic testing was still in its infancy, and the business case for the Indian market, where Tariyal was hoping to sell her product, was hard to make to American venture capitalists.

In 2011 she went to work for Pardis Sabeti, a Harvard professor who needed someone to manage a large genetic study in West Africa. It was in Sabeti's lab that she met Stephen Gire. The two of them traveled through Sierra Leone together to collect samples from survivors of Lassa fever, a deadly hemorrhagic fever broadly similar to Ebola.

Then, in 2013, Tariyal received a fellowship at Harvard Business School designed to encourage graduates to start new life-sciences companies. She was 33 at the time and an aspiring entrepreneur. She was not ready to have children and asked her doctor if she could wait five more years before she tried. She wanted to do a blood test called an anti-Müllerian hormone (or AMH) test that would approximate the number of viable eggs she had. But her doctor didn't see the need and wouldn't order it for her. And she was shocked by what the doctor suggested as an alternative: simply try to get pregnant to find out if she could.

This left Tariyal so alienated and frustrated that she decided her only option was to create her own AMH test that women could perform themselves, at home. She called Gire to ask for his help. She wanted to design assays to pick up proteins that would let her determine whether AMH and other hormones could be detected in menstrual blood, instead of blood drawn from veins, so that you wouldn't have to see a doctor to get tested. A woman could, in theory, just send in a used tampon for analysis.

During her fellowship, Tariyal performed tests that looked at three types of samples—venous blood, blood from a pinprick to the skin, and menstrual blood—to see where they overlapped. "I literally had to run them to a lab to process right away," she recalls. She was putting the logistical prowess she'd honed in Sierra Leone to use. As a menstruating woman, Tariyal also had an advantage: not only could she include herself in trials, but she was entitled to look at her own results.

To her disappointment, she found that AMH levels are consistently lower in menstrual blood than they are in venous blood. Her initial idea wouldn't work. But she believed she'd stumbled onto something even better: clear genomic signals in menstrual blood. Though genomics hadn't been her goal, it was a field rich with possibility. She found some 800 genes that were expressed differently in menstrual effluence and venous blood. The effluence contains not only blood but also endometrial lining, and some cervical and vaginal cells as well. It is, she says, like "getting a natural biopsy from your body."

With funding of $100,000 and six months of access to genome sequencing equipment at a startup accelerator run by the genomics company Illumina, she and Gire continued to look at menstrual blood samples. In particular, they hoped they might be able to reliably detect changes in gene expression that Linda Giudice, a doctor at the University of California, San Francisco, had recently discovered in women with endometriosis.

They have yet to succeed. Diagnosing diseases from menstrual blood is difficult. Published data establishing the efficacy of such diagnoses remains sparse, though sequencing technologies and other methods of extracting information from blood samples are fast improving. But NextGen Jane's access to the Illumina equipment ran out in 2015 (although it now uses equipment shared by a collective of genomics companies).

CHRISTIE HEMM KLOK

It is, she says, like “getting a natural biopsy from your body.”

BRUCE PETERSON

THE “WOMEN’S HEALTH” STIGMA

NextGen Jane is part of a cluster of firms trying to develop direct-to-consumer tests for endometriosis and other diseases affecting women.

As with any such boom, the surge of femtech companies leaves plenty to be wary of. The rise and fall of Theranos, which falsely claimed to have developed a revolutionary new method of blood analysis, has made people suspicious of biotech startups claiming to reinvent the blood test. A 2016 study by Columbia University researchers found that the overwhelming majority of menstrual tracking apps were inaccurate. Some defaulted to 28-day cycle lengths, though fewer than 15% of women have cycles precisely that long. Other apps predict a baby’s gender from the date of conception, or peddle other pseudoscientific claims.

Tariyal ultimately hopes to use menstrual blood to screen not only for endometriosis but also for cervical cancer and various other disorders. NextGen Jane’s key patent, at the moment, is for a device that wrings blood out of tampons. I watched her manipulate it. She seals a container and twists the mechanism like a pepper shaker. It squeezes out the blood into a compartment below.

The device has yet to be approved by the US Food and Drug Administration, but Tariyal says a clinical trial is designed and ready to go. She says she needs to raise several million more dollars to run a trial on about 800 women that could establish the diagnostic efficacy of menstrual blood. It will take her about two years, she says—if she can raise the money.

In a Washington Post op-ed last year, Tariyal outlined some of the challenges in fund-raising for a women’s health startup. “Some of my mentors recommended I mask the technology itself: Strip the deck of ‘menstrual blood’ and call it a novel female substrate, they suggested. Don’t say you’re a ‘women’s health’ company. It signals a lack of scientific heft,” she wrote. “I understood them to mean: Try to look as little as possible like what you really are—a woman-led company utilizing female biology to advance health care for half the population.”

Dayna Evans is a freelance writer based in New York.

Can AI pass the smell test?

Deploying artificial intelligence can be tricky and expensive. Companies had better know why they really want it.

BY Brian Bergstein

ILLUSTRATION BY Derek Brahney

The art of making perfumes and colognes hasn't changed much since the 1880s, when synthetic ingredients began to be used. Expert fragrance creators tinker with combinations of chemicals in hopes of producing compelling new scents. So Achim Daub, an executive at one of the world's biggest makers of fragrances, Symrise, wondered what would happen if he injected artificial intelligence into the process. Would a machine suggest appealing formulas that a human might not think to try?

Daub hired IBM to design a computer system that would pore over massive amounts of information—the formulas of existing fragrances, consumer data, regulatory information, on and on—and then suggest new formulations for particular markets. The system is called Philyra, after the Greek goddess of fragrance. Evocative name aside, it can't smell a thing, so it can't replace human perfumers. But it gives them a head start on creating something novel.

Daub is pleased with progress so far. Two fragrances aimed at young customers in Brazil are due to go on sale there in June. Only a few of the company's 70 fragrance designers have been using the system, but Daub expects to eventually roll it out to all of them.

However, he's careful to point out that getting this far took nearly two years—and it required investments that still will take a while to recoup. Philyra's initial suggestions were horrible: it kept suggesting shampoo recipes. After all, it looked at sales data, and shampoo far outsells perfume and cologne. Getting it on track took a lot of training by Symrise's perfumers. Plus, the company is still wrestling with costly IT upgrades that have been necessary to pump data into Philyra from disparate record-keeping systems while keeping some of the information confidential from the perfumers themselves. "It's kind of a steep learning curve," Daub says. "We are nowhere near having AI firmly and completely established in our enterprise system."

The perfume business is hardly the only one to adopt machine learning without seeing rapid change. Despite what you might hear about AI sweeping the world, people in a wide range of industries say the technology is tricky to deploy. It can be costly. And the initial payoff is often modest.

It's one thing to see breakthroughs in artificial intelligence that can outplay grandmasters of Go, or even to have devices that turn on music at your command. It's another thing to use AI to make more than incremental changes in businesses that aren't inherently digital.

AI might eventually transform the economy—by making new products and new business models possible, by predicting things humans couldn't have foreseen, and by relieving employees of drudgery. But that could take longer than hoped or feared, depending on where you sit. Most companies aren't generating substantially more output from the hours their employees are putting in. Such productivity gains are largest at the biggest and richest companies, which can afford to spend heavily on the talent and technology infrastructure necessary to make AI work well.

This doesn't necessarily mean that AI is overhyped. It's just that when it comes to reshaping how business gets done, pattern-recognition algorithms are a small part of what matters. Far more important are organizational elements that ripple from the IT department all the way to the front lines of a business. Pretty much everyone has to be attuned to how AI works and where its blind spots are, especially the people who will be expected to trust its judgments. All this requires not just money but also patience, meticulousness, and other quintessentially human skills that too often are in short supply.

Looking for unicorns

Last September, a data scientist named Peter Skomoroch tweeted: "As a rule of thumb, you can expect the transition of your enterprise company to machine learning will be about 100x harder than your transition to mobile." It had the ring of a joke, but Skomoroch wasn't kidding. Several people told him they were relieved to hear that their companies weren't alone in their struggles. "I think there's a lot of pain out there—inflated expectations," says Skomoroch, who is CEO of SkipFlag, a business that says it can turn a company's internal communications into a knowledge base for employees. "AI and machine learning are seen as magic fairy dust."

Among the biggest obstacles is getting disparate record-keeping systems to talk to each other. That's a problem Richard Zane has encountered as the chief innovation officer at UC Health, a network of hospitals and medical clinics in Colorado, Wyoming, and Nebraska. It recently rolled out a conversational software agent called Livi, which uses natural-language technology from a startup called Avaamo to assist patients who call UC Health or use the website. Livi directs them to renew their prescriptions, books and confirms their appointments, and shows them information about their conditions.

Zane is pleased that with Livi handling routine queries, UC Health's staff can spend more time helping patients with complicated issues. But he acknowledges that this virtual assistant does little of what AI might eventually do in his organization. "It's just the tip of the iceberg, or whatever the positive version of that is," Zane says. It took a year and a half to deploy Livi, largely because of the IT headaches involved with linking the software to patient medical records, insurance-billing data, and other hospital systems.

Similar setups bedevil other industries, too. Some big retailers, for instance, save supply-chain records and consumer transactions in separate systems, neither of which is connected to broader data storehouses. If companies don't stop and build connections between such systems, then machine learning will work on just some of their data. That explains why the most common uses of AI so far have involved business processes that are siloed but nonetheless have abundant data, such as computer security or fraud detection at banks.

Even if a company gets data flowing from many sources, it takes lots of experimentation and oversight to be sure that the information is accurate and meaningful. When Genpact, an IT services company, helps businesses launch what they consider AI projects, "10% of the work is AI," says Sanjay Srivastava, the chief digital officer. "Ninety percent of the work is actually data extraction, cleansing, normalizing, wrangling."

This doesn't mean AI is overhyped. But algorithms are a small part of what really matters in reshaping how business gets done.

Those steps might look seamless for Google, Netflix, Amazon, or Facebook. But those companies exist to capture and use digital data. They're also luxuriously staffed with PhDs in data science, computer science, and related fields. "That's different than the rank and file of most enterprise companies," Skomoroch says.

Indeed, smaller companies often require employees to delve into several technical domains, says

Anna Drummond, a data scientist at Sanchez Oil and Gas, an energy company based in Houston. Sanchez recently began streaming and analyzing production data from wells in real time. It didn't build the capability from scratch: it bought the software from a company called MapR. But Drummond and her colleagues still had to ensure that data from the field was in formats a computer could parse. Drummond's team also got involved in designing the software that would feed information to engineers' screens. People adept at all those things are "not easy to find," she says. "It's like unicorns, basically. That's what's slowing down AI or machine-learning adoption."

Fluor, a huge engineering company, spent about four years working with IBM to develop an artificial-intelligence system to monitor massive construction projects that can cost billions of dollars and involve thousands of workers. The system inhales both numeric and natural-language data and alerts Fluor's project managers about problems that might later cause delays or cost overruns.

Data scientists at IBM and Fluor didn't need long to mock up algorithms the system would use, says Leslie Lindgren, Fluor's vice president of information management. What took much more time was refining the technology with the close participation of Fluor employees who would use the system. In order for them to trust its judgments, they needed to have input into how it would work, and they had to carefully validate its results, Lindgren says.

To develop a system like this, "you have to bring your domain experts from the business—I mean your best people," she says. "That means you have to pull them off other things." Using top people was essential, she adds, because building the AI engine was "too important, too long, and too expensive" for them to do otherwise.

The seeds of AI

Once an innovation arises, how quickly will it diffuse through the economy? Economist Zvi Griliches came up with some fundamental answers in the 1950s—by looking at corn.

Griliches examined the rates at which corn farmers in various parts of the country switched to hybrid varieties that had much higher yields. What interested him was not so much the corn itself but the value of hybrids as what we would today call a platform for future innovations. "Hybrid corn was the invention of a method of inventing, a method of breeding superior corn for specific localities," Griliches wrote in a landmark paper in 1957.

Hybrids were introduced in Iowa in the late 1920s and early 1930s. By 1940 they accounted for nearly all corn planted in the state. But the adoption curve was nowhere near as steep in places like Texas and Alabama, where hybrids were introduced later and covered about half of corn acreage in the early 1950s. One big reason is that hybrid seeds were more expensive than conventional seeds, and farmers had to buy new ones every year. Switching to the new technology was a riskier proposition for the farms in these states than in the richer and more productive corn belt of the Midwest.

What Griliches captured, and what subsequent economists confirmed, is that the spread of technologies is shaped less by the intrinsic qualities of the innovations than by the economic situations of the users. The users' key question is not, as it is for technologists, "What can the technology do?" but "How much will we benefit from investing in it?"

Machine learning is making Facebook, Google, and Amazon rich. But outside that AI belt, things are moving much more slowly.

Today machine learning is undergirding every aspect of the operations of companies like Facebook, Google, and Amazon and many startups. It's making these companies exceptionally rich. But outside that AI belt, things are moving much more slowly, for rational economic reasons.

At Symrise, Daub thinks the perfume AI project fell into a sweet spot. It was a relatively small-scale experiment, but it involved real work for a fragrance client and wasn't a mere lab simulation.

"We're all under a lot of pressure," he says. "No one really has time to do greenfield learning on the side." Yet even this required a leap of faith in the technology. "It's all about conviction," he says. "There's a very strong conviction in me that AI will play a role in most of the industries we see today, some more predominantly. To completely ignore it is not an option."

Brian Bergstein is editor at large of Neo.life and a former editor at MIT Technology Review.

The 10 worst technologies of the 21st century

We all make mistakes sometimes.

By the editors
Illustrations by Daniel Savage

You'd think it would be easy to come up with a list of bad technologies from the past couple of decades. But we had a hard time agreeing: What makes a "bad" technology?

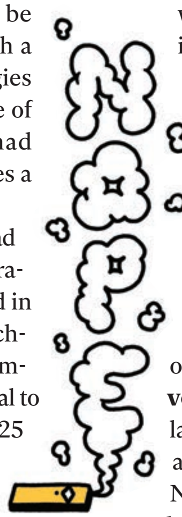

After all, technologies can be bad because they fail to achieve admirable aims, or because they succeed in wicked ones. The most useful technologies can also be the most harmful—think of cars, which are crucial to the modern world yet kill over 1.25 million people a year. And when well-intentioned technologies fail, is it because they are fundamentally flawed or just ahead of their time?

Take the **Segway**. Inventor Dean Kamen hyped it as a device that would transform cities and transportation. It turned out to be an expensive scooter that makes you look silly. Hoverboards were similarly all the rage until their batteries started exploding. But now (smaller) scooters and (safer) powered skateboards are increasingly popular.

If **Google Glass** had been developed by a lesser company, we probably wouldn't pick on it so much. But Google should have known better. It made the wearer appear elitist and invasive. Then again, like Segways and hoverboards, this was a failed product, not a failed technology; augmented-reality glasses and heads-up displays are finding their public.

Some technologies are well-intentioned but solve no real problems and create new ones. Before **electronic voting**, automated tabulating of paper ballots left an auditable paper trail. Now elections are more vulnerable to hacking.

Some failures apply a technological fix to what is really a social or political problem. Take **One Laptop per Child**, which set out to solve inequality in education with a new gadget. But was it simply too early? Commercial laptops, tablets, and—above all—smartphones have since inundated the developing world.

Indiscriminate uses of technology worry us. Sometimes this is because regulations are flouted. Gene-editing techniques like CRISPR may one day cure all manner of diseases, but right now we don't know if CRISPR is safe to use in humans. That's why the **CRISPR babies** born in 2018 make our list.

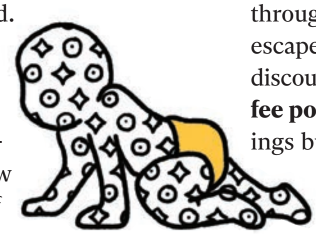

Other times, it's because technology has outpaced regulation. **Data trafficking**, the sharing and remixing of people's data without their control or awareness, has contributed to the undermining of personal liberty and democracy itself.

Some technologies are just misapplied. So far **cryptocurrency** looks mainly like a way for a handful of speculators to get very rich while a lot of other people end up poorer. But the technology underlying it, blockchain, could yet be transformative in other areas.

Still, there are a few inventions we could agree have no redeeming features. Juul and other **e-cigarettes** are addicting a new generation to nicotine, through a loophole that allowed them to escape public health regulations meant to discourage cigarette smoking. **Plastic coffee pods** save half a minute in the mornings but produce tons of hard-to-recycle waste. And as for **selfie sticks** … need we say more?

MIT Technology Review (ISSN 1099-274X), March/April 2019 issue, Reg. US Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2019. 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.


FAST AS h
THE LEXUS HYBRID LINE
Climb behind the wheel. Strap yourself in. Bury the pedal. Feel the roar. Then tighten your grip as the LC 500h's lightning-fast Multistage Hybrid Drive system propels you from zero to 60 in a mere 4.7 seconds.[1,2] Lexus Hybrids. There's more to h than just hybrid.
INSTANT ACCELERATION[3]
BOLD STYLING
COMPARABLY PRICED TO GAS MODELS[3]
LEXUS
EXPERIENCE AMAZING
LC 500h
Options shown. 1. Ratings achieved using the required premium unleaded gasoline with an octane rating of 91 or higher. If premium fuel is not used, performance will decrease. 2. Performance figures are for comparison only and were obtained with prototype vehicles by professional drivers using special safety equipment and procedures. Do not attempt. 3. Comparing 2018 Lexus Gas models and 2018 Lexus Hybrid models. Excluding special orders on the 2018 GS Hybrid. ©2018 Lexus