MIT Technology Review

VOLUME 117, NUMBER 4 | JULY/AUGUST 2014 | $5.99

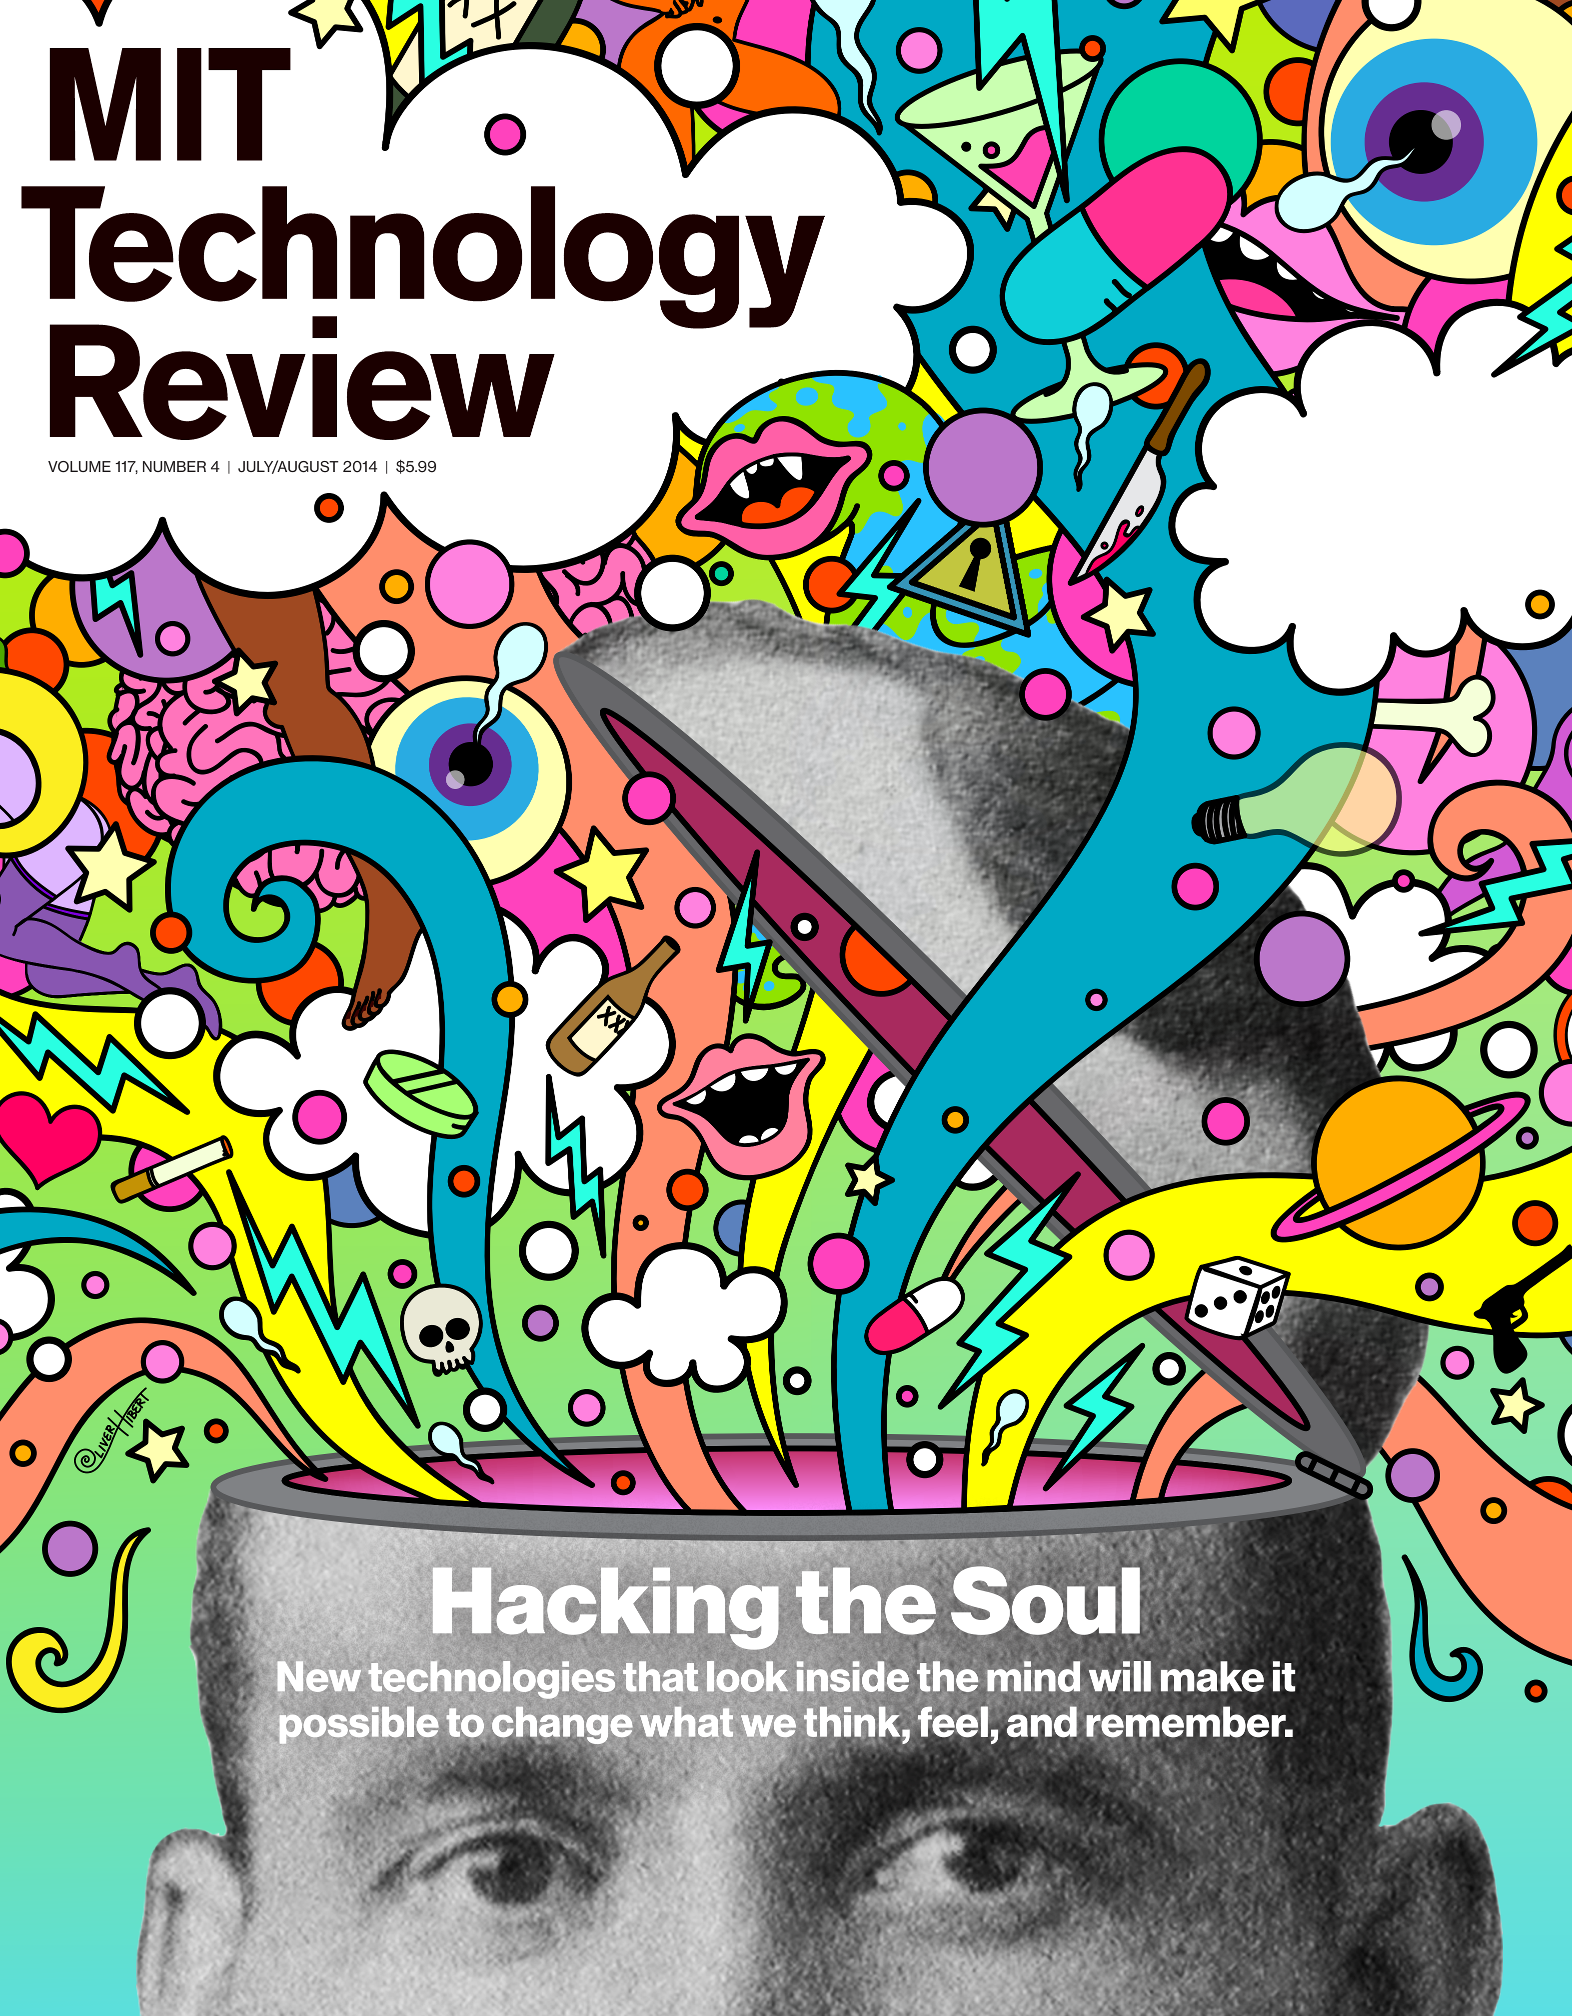

Hacking the Soul

New technologies that look inside the mind will make it possible to change what we think, feel, and remember.

So quiet, it's left our competition speechless.

K900: Certified to deliver a quieter cabin and a smoother, better overall ride than the Audi A8 4.0T, BMW 750i and Lexus LS 460.*

The Kia K900 was specifically designed to elevate the driving experience to a new level of luxury — one that even our most esteemed competitors have yet to achieve.

With a high-strength tensile steel structure and multi-link, fully independent front and rear suspension, the K900 delivers one of the smoothest rides available. It also silences the road with a foam-padded, sound-insulated cabin and laminated window glass. Because sometimes the most luxurious sound is nothing at all. Challenge the luxury you know.™

K900 V8 expected Spring 2014. Initially available in certain markets with limited availability. K900 V8 with VIP package shown. Not all optional features are available on all trims. Some features may vary. *AMCI Testing Certified, comparably equipped MY 2014 vehicles; K900 VIP trim, Lexus LS 460 with optional air suspension.


kia.com/K900
K900 V8 with VIP package shown.

From the Editor

LUDWIG WITTGENSTEIN WAS NOT TECHnologically illiterate: he had studied aeronautical engineering at the University of Manchester before World War I, researching the behavior of kites and designing a propeller with little rockets at its tips. But the philosopher was violently opposed to *scientism*, which Trenton Jerde, in his review of James Klagge's *Wittgenstein in Exile,* describes as "a preoccupation with the scientific method, the appeal to the sciences to solve problems that are beyond their reach, and a misuse of scientific terminology." The philosopher insisted on an unbridgeable divide between philosophy and science, which Klagge calls "Wittgenstein's insulation thesis," one of whose consequences was that science cannot resolve philosophical problems.

Wittgenstein would have been especially derisive of the claims of neuroscientists to meaningfully explain mental phenomena. In his *Zettel* (or posthumously collected remarks), he writes about the psychology of his own time: "No supposition seems to me more natural than that there is no process in the brain correlated … with thinking; so that it would be impossible to read off thought processes from brain processes."

Wittgenstein's interdiction is now a commonplace among philosophers. Many argue that understanding the causes of events in our brains cannot tell us much about the mind, because inferring anything about the latter from the former is a kind of "category mistake." But the attitude is becoming a rearguard defense against the encroachments of an advancing explanatory method. Questions such as "What is consciousness?" or "Do we have free will?" or "How do we ethically reason?" are of abiding interest, and because philosophers have made little progress in answering them, neuroscientists have felt at liberty to try. Thinking, feeling, and deciding are the most intimately human of all things, and yet we understand them hardly at all.

That neuroscientists can make the attempt is the result of recent technological advances, including (but not limited to) new kinds of brain imaging and the emerging field of optogenetics. This issue of *MIT Technology Review* describes those emerging technologies (see "Neuroscience's New Toolbox," by Stephen S. Hall, page 20, and "Cracking the Brain's Codes," by Christof Koch and Gary Marcus, page 42) and explains some of the surprising early insights they have suggested (see "Searching for the 'Free Will' Neuron," by David Talbot," page 64, and interviews with the neuroscientists Joseph LeDoux on memory, page 30; Antonio Damasio on emotions, page 48; and Rebecca Saxe on empathy, page 60). Finally, we report on the interventions the new technologies may make possible, including treatments for intractable mental illnesses such as schizophrenia (see "Shining Light on Madness," by David Rotman, page 34) and the use of brain-machine interfaces to help paralyzed patients ("The Thought Experiment," by Antonio Regalado, page 52).

Are the philosophers convinced by any of this? Not really. Responding to Gabriel Kreiman's research into decision-making, Hilary Bok, a philosopher at Johns Hopkins, is reserved: "I love these experiments and think they are really interesting, but I'm less convinced whether they have shown anything crucial about free will." But they are intrigued. Patricia Churchland, a philosopher at the University of California, San Diego, says of the same experiments, "Self-control is an entirely real brain phenomenon. Insofar as self-control is a key component of free choice, we do in fact have free choice."

But perhaps it doesn't matter much what professional philosophers think. They've had 2,000 years to answer these questions in their own way. The power of an explanation is its capacity to satisfyingly illuminate something hitherto obscure and to allow us to do things we could not before (here, effectively treat mental illnesses and build brain prosthetics). Insofar as traditional philosophy has an important role in understanding the mind, it may be to pose questions and parse answers, and the questions we ask will become more interesting because of the conceptual breakthroughs of neuroscience.

But write to me at jason.pontin@technologyreview.com and tell me what you think.

GUIDO VITTI


CASE #31345
HU848
31266
WHEN YOU ARE PASSIONATE ABOUT SOMETHING,
YOU PROTECT IT IN THE BEST WAY.
ZURICH INSURANCE.
FOR THOSE WHO TRULY LOVE THEIR BUSINESS.
ZURICH®

Contents

Front

July/August 2014

Back

ON THE COVER

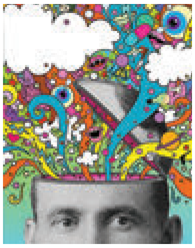

Illustration by Oliver Hibert for MIT Technology Review

SCULPTURE: JOSHUA HARKER; PHOTOGRAPH: BRUCE PETERSON

Editor in Chief and Publisher
Jason Pontin

EDITORIAL

Editor
David Rotman

Deputy Editor
Brian Bergstein

Creative Director
Eric Mongeon

News and Analysis Editor
Will Knight

Chief Correspondent
David Talbot

Senior Editor, Energy
Kevin Bullis

Senior Editor, Business Reports
Nanette Byrnes

Senior Editor, MIT News
Alice Dragoon

Senior Editor, Business
Antonio Regalado

Senior Editor, IT
Tom Simonite

Senior Web Producer
Kyanna Sutton

Managing Editor
Timothy Maher

Copy Chief
Linda Lowenthal

IT Editor, Web and Social Media
Rachel Metz

Biomedicine Editor
Susan Young Rojahn

Research Editor
Mike Orcutt

Assistant Art Director
Colin Jaworski

Editorial Assistant
J. Juniper Friedman

Production Director
James LaBelle

Contributing Editors
Katherine Bourzac
Jon Cohen
Peter Fairley
Simson L. Garfinkel
Robert D. Hof
Courtney Humphries
Martin LaMonica
John Pollock
Amanda Schaffer

CORPORATE

Chief Financial Officer
Rick Crowley

Chief Operating Officer
James Coyle

Chief Strategy Officer
Kathleen Kennedy

Director of International Business Development
Antoinette Matthews

Executive Assistant
Leila Snyder

Manager of Information Technology
Colby Wheeler

Office Manager
Linda Cardinal

PRODUCT DEVELOPMENT

Chief Digital Officer and VP, Product Development
Erik Pelletier

Project Manager
Jane Allen

Senior Software Engineers
Shaun Calhoun
Molly Frey

User Interface/Digital Designer
Emily Dunkle

EVENTS

Director of Events and Strategic Partnerships
Amy Lammers

Director of Events Programming
Laura Janes Wilson

Senior Content Producer
Marcy Rizzo

Event Coördinator
Lina Umansky

CONSUMER MARKETING

VP, Consumer Revenues and Marketing
Bruce Rhodes

Marketing Communications Manager
David W.M. Sweeney

FINANCE

General Ledger Manager
Olivia Male

Accountant
Letitia Trecartin

ADVERTISING SALES

Director of Advertising Sales
James Friedman

Midwest Sales Director
Maureen Elmaleh
maureen.elmaleh@technologyreview.com
303-975-6381

New England, Detroit, and Canada
Barry Echavarria
barry.echavarria@technologyreview.com
603-924-7586

Mid-Atlantic and Southeast
Clive Bullard
cbullards@cs.com
845-231-0846

West Coast
Rob Finley
rob.finley@technologyreview.com
415-659-2982

Europe
Anthony Fitzgerald
mail@afitzgerald.co.uk
44-1488-680623

France
Philippe Marquezy
philippe.marquezy@espacequadri.com
33-1-4270-0008

Germany
Michael Hanke
michael.hanke@heise.de
49-511-5352-167

China
Tao Lin
imlintao@hotmail.com

Japan
Akiyoshi Ojima
ojima@media-jac.co.jp
813-3261-4591

Spain and South America (Online)
Pedro Moneo Laín
pedro.moneo@opinno.com
34-667-732-894

Advertising Services Coördinator
Ken Collina

Sales & Marketing Associate
Julie Swanson

Custom Editor
Anne Stuart

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

Media Kit
www.technologyreview.com/media

MIT ENTERPRISE FORUM, INC.

Executive Director
Antoinette Matthews

Senior Producer, Content and Community
Susan Kaup

Director of Chapter Leadership and Process
Gaylee Duncan

Director of Communications
Joyce Chen

Chairman
Jason Pontin

President
Kathleen Kennedy

Treasurer
James Coyle

CUSTOMER SERVICE AND SUBSCRIPTION INQUIRIES

National: 800-877-5230

International: 818-487-2088

E-mail: technologyreview@pubservice.com

Web: www.technologyreview.com/customerservice

MIT Records: 617-253-8270 (alums only)

Reprints: Donna Summers
dsummers@wrightsmedia.com
281-419-5725

Licensing and permissions:
licensing@technologyreview.com

Technology Review
One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

MIT Technology Review identifies the new technologies that matter—explaining how they work, deciphering their impact, and revealing how they'll change our lives. Founded at the Massachusetts Institute of Technology in 1899, we are a digitally oriented global media company that publishes serious journalism in magazines and on websites and mobile devices. We also host events around the world. 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.

De Technologia non multum scimus. Scimus autem, quid nobis placeat.

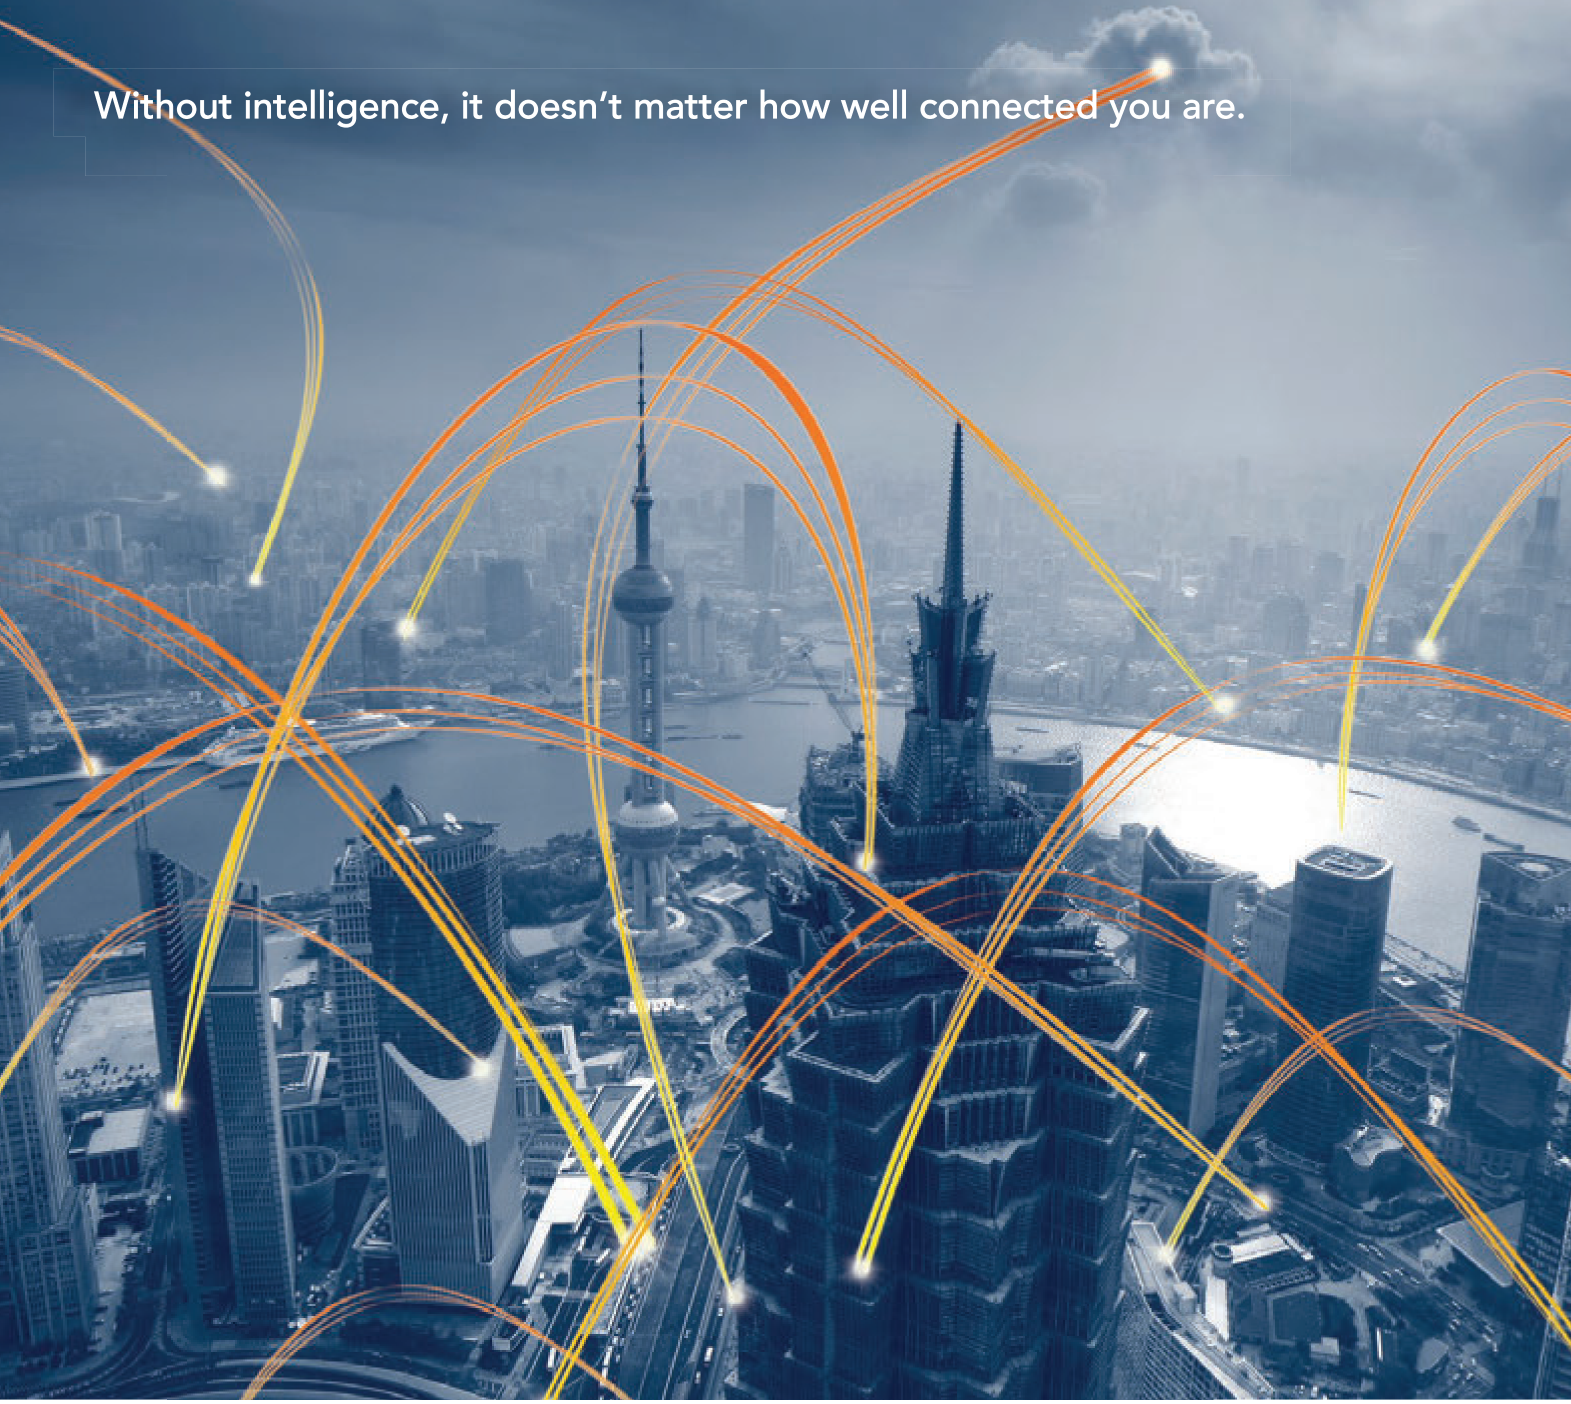

Bringing intelligence to the Internet of Things

Cars communicating with highways to avoid accidents and traffic congestion. Home hubs talking to the smart grid to reduce energy consumption. Medical monitoring from across town. The billions of intelligent connections that make our world smarter, greener and safer. This is Freescale Connected Intelligence. Make it work for you. Visit **freescale.com/intelligence**

make it
freescale™

Feedback

E-mail letters@technologyreview.com

Write *MIT Technology Review*
One Main Street, 13th Floor
Cambridge, MA 02142

Please include your address, telephone number, and e-mail address.

Letters and comments may be edited for both clarity and length.

Five Most Popular Stories

MIT Technology Review
Volume 117, Number 3

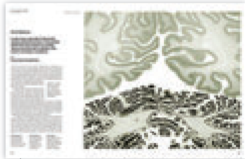

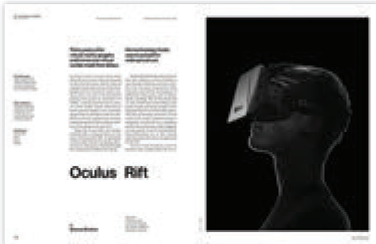

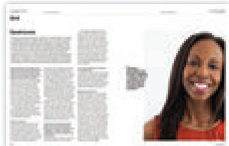

1 Smart Wind and Solar

If the grid can't manage the tiny fraction of power now produced by wind and solar, just imagine how bad the problem would be if they were increased five times or more. **—big.league.slider**

@big.league.slider: It's actually easy to imagine, just like in Colorado or Texas when inclement weather throws thermal and nuclear plants offline and poor old intermittent power has helped utilities to ride it through. That's what this article implies: stronger modeling allows utilities to understand better what they are facing. It's a learning curve. **—jeffhre**

2 Ultraprivate Smartphones

What if Silent Circle and Blackphone are actually operating by order of the NSA, set up to catch exactly those who really value privacy and therefore are prime targets? Do we have any guarantees? **—bluppfisk**

Sometimes I wonder about this fanatical need for secrecy. I take my wife to an Indian restaurant and the waiters begin preparing her favorite samosas when we drive into the parking lot. It is part of life and a nice part of it at that. I don't want to walk around being suspicious of everyone and everything in my vicinity. **—tag**

3 Oculus Rift

Anybody who has had the opportunity to try side by side the Google Glass and the Oculus Rift (as I could at a recent conference in Brussels) would quickly understand Zuckerberg's acquisition decision. Between the "cool" factor and the "wow!" effect, the fact that Glass is less obstructive (allegedly allowing you to keep in contact with the real world) may eventually play against Glass commercially. Young people might be more reluctant to look like a "dandy geek" (as they would with Glass) than to stake their claim as a "dedicated nerd" (as they would with Oculus Rift). **—borisazais**

4 Brain Mapping

We will never mimic the human brain in silicon or carbon. The physiology of neurons is not nearly as intriguing as the behavior that allows them to self-assemble into pattern-matching engines that can be recalled and reapplied. **—marcvky**

It would have been nice if you'd addressed what we expect to gain by seeing the brain cells. Can we decipher thoughts, ideas, fears, by just seeing? Can we understand conditions like autism? Can we enhance the brain functions of healthy individuals? What is the problem we are trying to solve, and how far are we from that? **—ankur**

5 Q&A: Sarah Lewis

Art is exploration and curiosity, so in that sense it shouldn't be considered much different from science. The scientific process—hypothesis, designed experiment, results analysis, and all over again—can be applied everywhere. **—vnedovic**

My art skills are minimal, but I've been interested in the relationship between art and science since college. Having worked in video-game design since the 1970s, I've seen brilliant engineers enjoy juggling, chemists coding intricate game logic, a whale-study scientist creating a game, and software engineers writing science fiction. **—ricklev**

More Nukes Is Not the Answer

In "Irrational Fears" (Views, May/June), Nathan Myhrvold tells us that the Fukushima accident followed numerous disastrous design decisions and was entirely preventable. The real lesson, he says, is that we should build modern reactors—but hindsight bias can be used rhetorically to trivialize design errors made 40 years ago. The mistakes made then weren't trivial or obvious at the time, and the mistakes being made today aren't, either. If they were, we wouldn't make them.

The lesson is that we need to try to reduce risk from critical design errors. Until we have more careful thinking on the pertinent issues, the simplest way to lower risks in nuclear reactor design is to have fewer reactors, not more.

Harold Thimbleby
Swansea University, Wales

Better Living Through Big Data?

Nicholas Carr's review of Sandy Pentland's book *Social Physics* ("The Limits of Social Engineering," May/June) reminded me of my work in the artificial-intelligence community in the 1980s. I think most of us working in the industry didn't believe the hype that surrounded AI at the time, but the hype was necessary to sustain the R&D dollars to keep even the more pragmatic developments going. Many of those developments, I might add, we see today in what Google, Apple, and Facebook have done. I personally enjoyed reading Pentland's book. We need the Pentlands of the world to show us the possibilities. —llocklee

Pentland is an engineer and perhaps overoptimistic that "engineering approximations" will allow him to do useful things with all that data. A computer scientist would probably take comfort in fundamental social phenomena that make societies unpredictable. —Dean L

Views

Philip Rosedale

Pamela Ronald

Vivek Wadhwa

COMPUTING

Virtually There

A new generation of virtual reality will compete with the physical world.

RECENT MONTHS HAVE BEEN GOOD ones for virtual reality. The Facebook acquisition of Oculus has galvanized the idea that "something wonderful" will happen if we put on these strange headsets and visually enter other worlds.

Most people assume this means gaming. And while it's true that the new headsets will immerse users in amazing gaming experiences, that's not the big part of the story.

After we've tried the Oculus for a few months and the novelty has worn off, we might find ourselves asking: "Where are the other people?" That is where things are going to get interesting.

The Oculus Rift is only one of several advances in hardware that are going to dramatically change our ability in the coming year to immerse ourselves in a 3-D world. For one thing, we'll be able to communicate with others while we're inside these worlds: the Internet is now fast enough to allow us to be in a virtual environment with other people who are accessing it from elsewhere, even halfway across the world.

Updating imagery shown to the eyes with a delay of less than 10 milliseconds relative to head movements generates a magical sense of being "present" in a virtual space. My own experiments have shown that a second kind of presence—the feeling of really being face to face with another person—requires an end-to-end delay (including hardware, software, and network transmission) of around 100 milliseconds or less between your movement and the other person's perception of that movement.

Below that threshold, the small head and eye movements that we use with each other while talking in the "real" world can work in a virtual one. We can feel empathy and connection, interrupt each other, and smoothly and rapidly exchange thoughts. At less than 100 milliseconds of delay you can reach out and virtually touch or shake hands with another person and find the perception of the resulting collisions and motion to be perfectly believable and immersive.

If virtual reality can replace (or even improve upon) videoconferencing or long-distance travel as a way of getting together with people, it will surely disrupt and restructure many basic human exchanges that have nothing to do with playing games.

For many of the everyday things we do—talking face to face, working together, or designing and building things—the real world will suddenly have real competition.

Philip Rosedale is CEO of High Fidelity, a startup working on a new virtual world. He was previously CEO of Second Life.

BIOTECHNOLOGY

Scare Tactics

How Vermont got it wrong on GMOs, and how that hurts people everywhere.

IN EARLY MAY, VERMONT GOVERNOR Peter Shumlin signed a bill into law that required a label for any foods produced with genetic engineering. This made Vermont the first U.S. state to require mandatory GMO labeling. (More than 50 countries already require such labels, and more than a dozen states are considering similar laws.) Many food activists viewed the bill's passage as a boon—GMOs cause cancer, they claim, and force farmers to use more toxic chemicals.

But if Vermont had honestly assessed genetically engineered crops, the bill

SAM KERR

would have indicated that there is not a single credible report of dangerous health effects from GMOs. It would have mentioned that the technology has been used safely in food and medicine for 30 years. It would have stated that farmers' use of GMO crops has *reduced* by a factor of 10 the amount of insecticide sprayed on corn over the last 15 years.

Aside from being disingenuous, the Vermont bill is inconsistent. It doesn't require labeling for cheese made with genetically engineered enzymes or red grapefruit developed through radiation mutagenesis. It doesn't require labeling for animals that have been fed GMO crops, or labels for crops sprayed with carcinogenic compounds. The law doesn't require crops sprayed with the organic pesticide Bt to be labeled, but crops genetically engineered to produce Bt must be labeled.

So the law won't give consumers access to food that's more sustainable, more healthful, or less "corporate."

Should you care? After all, any individual consumer can either pay heed to a label or ignore it. But political campaigns that reject science can have devastating consequences. Aside from the obvious example of climate skepticism, there's the antivaccination movement, which has led to outbreaks of measles and whooping cough.

The anti-GMO movement has already made consumers fearful. Many are now willing to pay more for "non-GMO" foods—providing an incentive for farmers to use older, more toxic, more expensive chemicals and farming practices.

GMO scaremongering campaigns have also harmed the poor. Vitamin A deficiency causes blindness in half a million people a year in the developing world. A genetically engineered rice called Golden Rice could have provided the nutrient at a fraction of the cost of current supplementation programs more than a decade ago. But regulatory bodies in India, Bangladesh, and elsewhere have still not approved Golden Rice for release. UC Berkeley agricultural economist David Zilberman calculates that swifter implementation would have saved at least one million people from blindness and prevented the death of thousands of children.

So let's label food, but let's do it right. Let's label food with details about how the crop was grown and what is actually in the food. Let's apply these labels to all foods, so consumers can draw their own conclusions about the risks and benefits. A right to know is only useful if the information is accurate.

Pamela Ronald is coauthor of Tomorrow's Table: Organic Farming, Genetics, and the Future of Food.

INTERNET

Our Lagging Laws

Making new laws is a slow process. Technology moves too fast to care.

LAWS FORBID LENDERS FROM discriminating on the basis of race, gender, and sexuality. Yet lenders can refuse to give a loan to people if their Facebook friends have bad payment histories, if their work histories on LinkedIn don't match their bios on Facebook, or if a computer algorithm judges them to be socially undesirable.

Such regulatory gaps exist because laws haven't kept up with technology. And the gaps are getting wider.

Technology now touches practically everyone, everywhere. Changes of a magnitude that once took centuries now happen in decades, sometimes in years. Not long ago, Facebook was a dorm-room dating site, mobile phones were for the ultrarich, drones were multimillion-dollar war machines, and supercomputers were for secret government research. Today, hobbyists build drones, and poor villagers in India access Facebook accounts on smartphones more powerful than the Cray 2 supercomputer, which in 1985 cost $17.5 million and weighed 2,500 kilograms.

We haven't come to grips with what is ethical in relation to technologies such as social media, let alone with what the laws should be. Consider the question of privacy. There is a public outcry today about surveillance by the National Security Agency, but the breadth of that surveillance pales in comparison with the data that Google, Apple, Facebook, and legions of app developers are collecting. Our smartphones track our movements and habits. Our Web searches reveal our thoughts. Where do we draw the line on what is legal—and ethical?

The Genetic Information Nondiscrimination Act of 2008 prohibits the use of genetic information in health insurance and employment. But it provides no protection from discrimination in long-term-care, disability, and life insurance. There are no laws to stop companies from using genomic data in the same way that lending companies and employers use social-media data. We will have similar debates about self-driving cars, drones, and robots.

As a society, we need to be mindful that powerful innovation now occurs too quickly for existing ethical frameworks and laws. This means questioning, rethinking, and reframing those values, as a culture, at an accelerated speed. And it means creating a legal system that can keep pace with a new era of techno-socioeconomic transformation.

Vivek Wadhwa is a fellow at Stanford University's Arthur and Toni Rembe Rock Center for Corporate Governance.

Upfront

PING ZHU

The Next Startup Craze: Food 2.0

Silicon Valley investors and startups are trying to improve our food. Do they bring anything to the table?

By Ted Greenwald

Most tech startups are silent spaces where earbud-wearing engineers peer into monitors. Not Hampton Creek Foods. The two-year-old company's office—a filled-to-bursting space in San Francisco's South of Market tech hotbed—grinds, clatters, and whirs like a laundromat run amok. That's the sound of industrial-strength mixers, grinders, and centrifuges churning out what the company hopes is a key ingredient in food 2.0: an animal-free replacement for the chicken egg.

Silicon Valley venture capitalists have funded several food-related startups in the past year, but Hampton Creek has gathered the most momentum. It has A-list investors including Founders Fund, Horizon Ventures, and Khosla Ventures, and two undisclosed industrial food companies are experimenting with its plant-based egg substitute. The prepared-food counter at Whole Foods began using the startup's egg-free Just Mayo in September 2013, and four other mainstream grocery chains are lined up for the first half of this year. And thanks to a recent investment round that boosted Hampton Creek's funding to $30 million and drew in Li Ka-shing, the wealthiest person in Asia, Just Mayo will soon be sold by a large online grocer in Hong Kong.

Hampton Creek Foods and other startups have big dreams of restructuring the food supply so that it uses less land, water, energy, and other resources. They are taking on corporate giants such as ConAgra, General Mills, and Kraft, which spend billions on research and technology development.

Similar ambitions have run up against considerable challenges in industries such as clean tech. But those involved in the food craze might prefer a different example. Hampton Creek's CEO, Josh Tetrick, wants to do to the $60 billion egg industry what Apple did to the CD business. "If we were starting from scratch, would we get eggs from birds crammed into cages so

Upfront

small they can't flap their wings, shitting all over each other, eating antibiotic-laden soy and corn to get them to lay 283 eggs per year?" asks Tetrick, a strapping former West Virginia University linebacker. While an egg farm uses large amounts of

water and burns 39 calories of energy for every calorie of food produced, Tetrick says he can make plant-based versions on a fraction of the water and only two calories of energy per calorie of food—free of cholesterol, saturated fat, and allergens, and without risk of avian flu or cruelty to animals. For half the price of an egg.

That's a tall order. The lowly chicken egg is a powerhouse of high-quality, low-cost nutrition. Yet in prepared foods, it's prized less for its nutrients than for its culinary properties: emulsifying, foaming, binding, gelling, and many more. Those functions are conferred by its unique complement of proteins. Rather than trying to reproduce the egg wholesale, Hampton Creek focuses on discovering vegetable proteins that replicate specific functions.

The rear of Hampton Creek's facility is devoted to finding them. There Josh Klein, a biochemist who formerly worked on an AIDS vaccine, runs a high-throughput screening pipeline designed to comb through millions of plant cultivars for proteins with certain characteristics. So far, Hampton Creek has examined 3,000 plants and discovered 11 desirable proteins, seven of which are already allowed in food by the U.S. Food and Drug Administration. One is the Canadian yellow pea protein that emulsifies the oil and water in Just Mayo. Another binds the company's cookie dough, called Eat the Dough, which is due to hit store shelves this month. A third goes into a prototype viscous yellow liquid that looks much like a beaten egg and possesses a similar nutritional profile, minus the cholesterol. When Hampton Creek chef Chris Jones, late of the Chicago restaurant Moto, pours the prototype into a heated pan, it solidifies as though it came out of a shell. The fake egg looks and feels authentic, but in a test last winter it was rife with off flavors, and the team is still working to make the formula satisfying under all possible conditions in which it might be prepared and eaten.

"I consider a broken food system my enemy," says Hampton Creek's CEO, Josh Tetrick.

Hampton Creek's emulsifiers and binders won't have the market to themselves. They'll compete with an array of existing egg substitutes—never mind eggs themselves. Egg replacers are widely used in packaged foods to cut costs and mitigate fluctuations in egg prices. These products, usually derived from soy, milk, gums, or starches, are typically cost-competitive with Hampton Creek's offering. Some of them also can be functionally more effective than eggs.

Competition isn't the only challenge. First, while Tetrick is intent on conquering the mainstream, his egg-free products risk being relegated to the vegan aisle. "I think their success is going to be limited to a niche market for now," says Christine Alvarado, an associate professor of poultry science at Texas A&M. And should the company discover a compelling protein that comes from an unusual plant, it must convince farmers to supply that crop in huge quantities without raising costs. "The more specialized your raw material, the higher the risks the supply chain will face," observes Jon Stratford of Natural Products, which makes a soy-based egg substitute for the food industry.

Then there's the fundamental question of whether Hampton Creek's proteins are, indeed, better than traditional egg replacements. In fact, Hampton Creek's own patent application offers egg replacement recipes made entirely of off-the-shelf ingredients (with additional processing such as grinding into very small particles). "Food scientists have been doing this for 100 years or more," says Gregory Ziegler, a professor of food science at Penn State University. (Tetrick says the recipes in the patent applications came from an earlier stage of research.)

Nonetheless, Hampton Creek's approach is working so far. Tetrick expects to have his egg-free mayo in 15,000 stores by late summer, up from 3,500 now, and he has his eye on fast-food chains and food service companies. He plans to triple the size of his workforce by the end of the year and expand his floor space thirtyfold. And once he has established a toehold in the egg market, Tetrick plans to start putting his plant-protein database to use in other areas; he hints at substitutes for chicken or beef. "I consider a broken food system my enemy," he says.

QUOTED

"OpenSSL is like a faulty engine part that's been used in every make and model of car, golf cart, and scooter."

—Jonathan Sander, of StealthBits, on a software library that left thousands of websites vulnerable to attack in April.

Talk of an Internet Fast Lane Is Hurting Some Entrepreneurs

The FCC's net neutrality plan will raise costs for companies that need fast connections or use a lot of bandwidth.

By David Talbot

Some venture capitalists at the cutting edge of Internet innovation say they will shun startups requiring fast connections for video, audio, or other services, mindful that ISPs may soon be able to charge extra fees to major content providers to ensure speedy service. Progress toward so-called Internet fast lanes has been accelerated by the U.S. Federal Communications Commission's recent vote to advance a plan that would allow such charges.

The cable industry says the charges are sensible, especially when a few content providers like Netflix can make up a large fraction of traffic. But if deep-pocketed players can pay for faster, more reliable service, then startups face a crushing disadvantage, says Brad Burnham, managing partner at Union Square Ventures, a VC firm based in New York City.

Burnham says that his firm will steer clear of startups working on video and media businesses, and it will also avoid investing in payment systems or in mobile wallets, which require ultrafast transaction times to make sense. "This is a bad scene for innovation in those areas," Burnham says.

History does suggest that some Web products have benefited from an open Internet. The founders of Foursquare, as an example, were able to set up their mobile social-networking service and reach 100,000 users with a mere $25,000 budget, Burnham says.

Other VCs, particularly those who fund broadband providers, say the explosion of video services has triggered massive costs that far exceed the growth in subscribers. Gillis Cashman, a managing partner at MC Partners in Boston, says it makes sense to charge extra to big content providers like Netflix, whose service at peak hours can account for more than 30 percent of total Internet traffic.

Less financially invested observers have little sympathy for this argument. Rob Faris, research director at the Berkman Center for Internet & Society at Harvard University, notes that broadband providers typically have had very high profit margins and often charge tiered rates based on the speeds consumers desire. "You can't credibly argue that consumers aren't paying enough for access to maintain higher-bandwidth services," he says.

If the FCC does let ISPs impose access fees, new business models and technologies for imposing those charges will emerge. Wireless carriers like Verizon are working on fast-lane technology, as are Web optimization companies like Akamai. And AT&T already offers so-called sponsored data that allows a broadcaster to subsidize its bits so you can watch shows (and ads) free on your smartphone, though other streams of data still count against your monthly caps.

TO MARKET

Personal Assistant

Oslo

COMPANY:
Microsoft

PRICE:
$9.99 per month (for Office 365)

AVAILABILITY:
Second half of 2014

Microsoft Oslo is a predictive virtual assistant for office workers. It draws on online content and a company's internal data to offer up important information, context, and contacts when they are needed, before you even think to ask. Oslo examines what you're working on to curate a selection of articles to read, Web pages to visit, videos to view, and podcasts, all presented on a page of clickable tiles. Currently available for testing, it will be included in Office 365, Microsoft's subscription productivity software for PCs and mobile devices.

COURTESY OF MICROSOFT (OSLO); COURTESY OF STEALTHBITS (SANDER)

Upfront

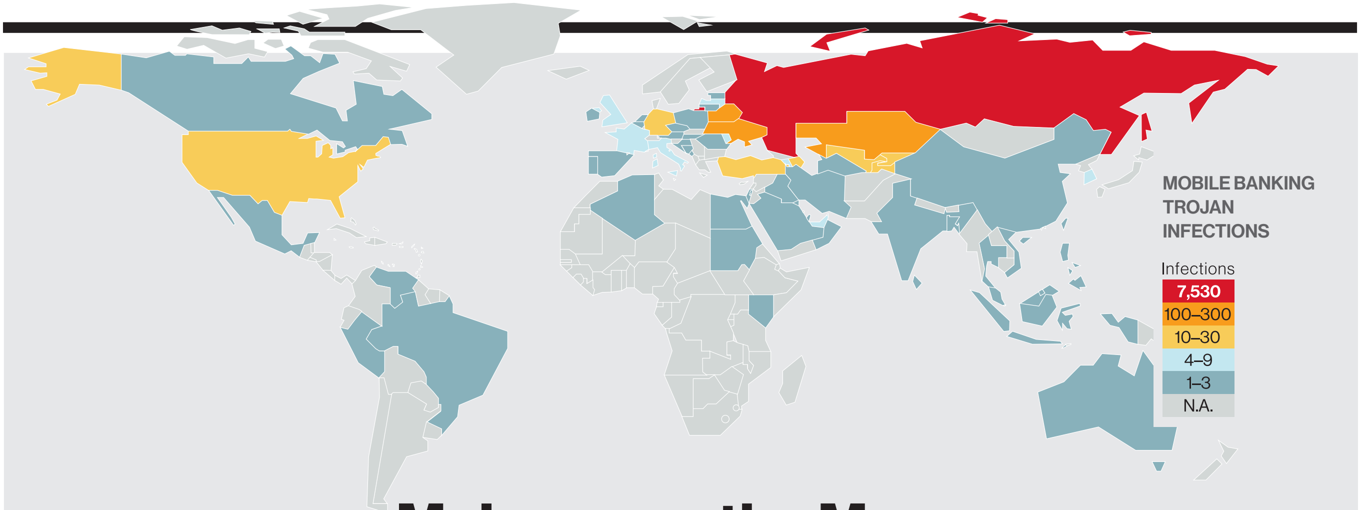

Malware on the Move

Mobile devices are increasingly used for financial transactions, and cybercriminals are taking notice. Trojan apps designed to steal log-in information are appearing much faster than other types of mobile malware, especially in Russia and other former Soviet states. The majority of malware is made for Google's Android operating system, which allows unscreened software to be installed, but as long as these schemes pay off, thieves may soon look to other platforms, in more parts of the world.

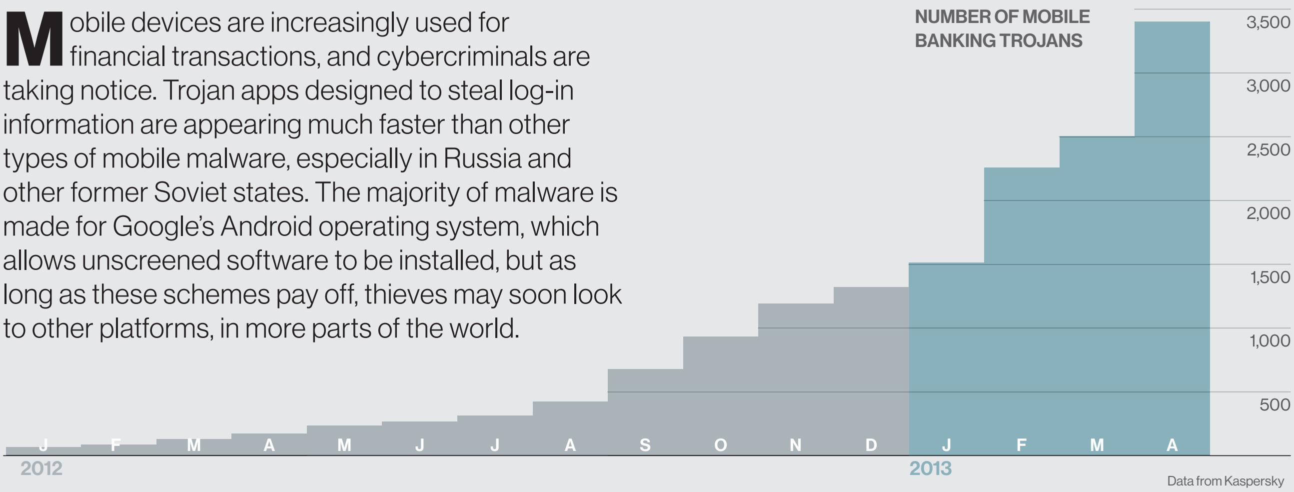

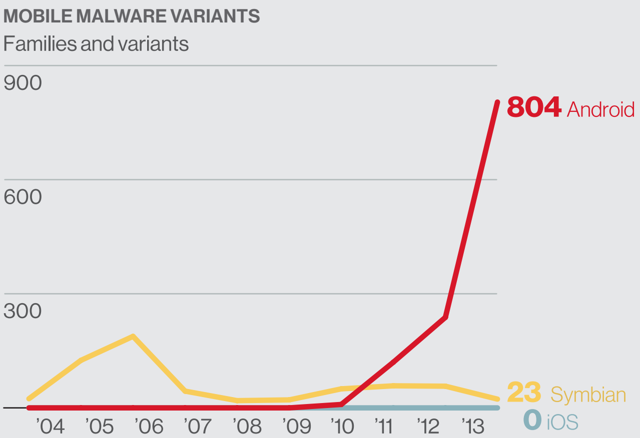

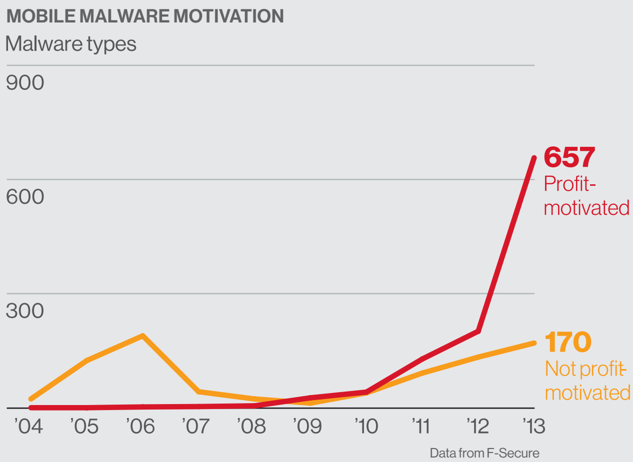

1.6 vs. 466
kilobits per second

Average Internet speed for users in the 20 slowest and fastest countries on earth.

Can Compounds in Young Blood Fix Aging?

Animal studies on the revitalizing power of young blood suggest new drug targets for treating age-related conditions.

By Susan Young Rojahn

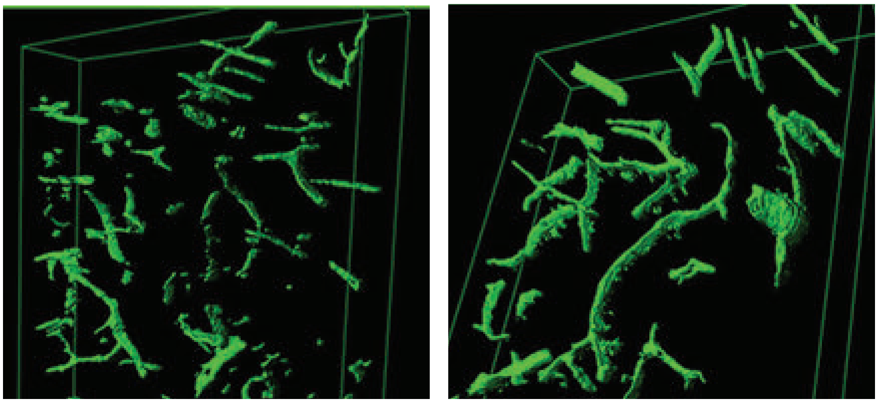

A 3-D reconstruction shows how the deterioration of blood vessels in the brains of older mice (left) can be reversed by treatment with young blood (right).

Researchers and investors are dreaming up medical treatments based on the near-fantastical finding that the blood of young mice can rejuvenate older mice. In some cases, a single protein found circulating in the blood is sufficient to restore muscle tissue and improve brain activity. The excitement is spurred by three studies published in *Science* in May, showing that components of blood from young mice, including a protein growth factor called GDF11, can repair damage and improve muscle and brain function in older mice.

The Boston-area venture capital firm Atlas Venture has started a still-unnamed company on the basis of the new findings and earlier results. Previous findings by Amy Wagers, a coauthor on two of the new studies, had caught the eye of the VC firm in 2013, and the new results "increased the excitement for the role of GDF11 in aging," says partner Peter Barrett. "Now it's the blocking and tackling of trying to understand what would be the best therapeutic approach to make this a commercial product."

The company might target heart failure first, since Wagers's team has shown that the growth factor can reverse age-related thickening of heart muscle in mice. GDF11 alone does not provide the full beneficial effects of young blood, says Wagers. "But now that we have a molecule that we can understand at a mechanistic level, we can start to build out the network of what is changing in the blood, find things that interact with GDF11, and regulate it to give us a clearer view of the network," she says.

Though the results are exciting, there are many outstanding questions. Potential side effects are not yet known. Other such growth factors, which encourage cell division, have been linked to cancer. Furthermore, blood-vessel growth is a common feature of cancer development, and genes involved in stem-cell function and regeneration have been linked to tumor growth. And researchers have never studied the treated mice for more than 60 days.

TO MARKET

Feel the Force

Touch

COMPANY:
3D Systems

PRICE:
$649

AVAILABILITY:
July

A new haptic stylus that lets you "feel" virtual objects and textures costs far less than existing devices, whose prices range from $2,400 to $13,000. It is part of a trend toward cheaper tools for 3-D imaging, modeling, and printing in a variety of applications. The product from which the Touch was derived, called Phantom, was originally invented by an MIT startup called Sensable, whose technology 3D Systems acquired. "There has always been a high cost barrier to doing this," says Ping Fu, chief entrepreneur officer at 3D Systems. "This brings that kind of ability right down into the consumer space, so anybody can pick up the haptic device and start sculpting."

COURTESY OF LIDA KATSIMPARDI (BLOOD); DAVID TALBOT/MIT TECHNOLOGY REVIEW (STYLUS)

Upfront

“Having information one minute earlier can actually be quite life-saving.”

—Steve Horng, an emergency physician at Beth Israel Deaconess Medical Center in Boston, who uses Google Glass to view patient records.

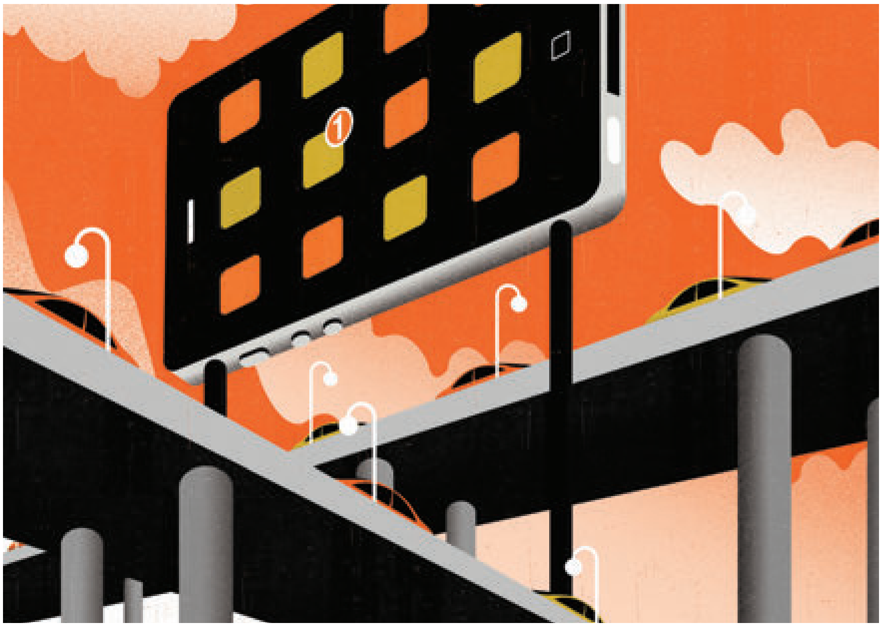

Startups Experiment with Ads That Know How You Drive

A mobile ad company plans to offer deals based on data collected from in-car apps.

By Tom Simonite

As businesses race to connect our homes and cars to the Internet, one mobile ad company scents a new opportunity. San Francisco–based Kiip plans to sell a new kind of ad targeted according to people’s actions at home or on the road, offering rewards or deals in exchange for certain behavior. “You get to your meeting early and you should get a free coffee from the place around the block,” Michael Sprague, head of partnerships for Kiip, said at the Ad:tech conference in San Francisco in March.

Kiip’s move comes at a time when more and more data on people’s actions is becoming available as wearable devices, Internet-connected home automation equipment, and cars with integrated data connections head to market. Those new data streams could form the basis for many new services and products, but they also bring new privacy concerns.

Ads tailored to driving behavior will be possible thanks to a partnership between Kiip and fellow startup Mojio. It expects to launch a $149 device this summer that plugs into a car’s diagnostic port and streams vehicle data to a smartphone app to help users track their driving, their fuel economy, and their vehicle’s maintenance status. Kiip will use data from that device to target promotions inside the Mojio app.

Sprague says that getting access to data from a car’s engine and safety systems could unlock new approaches to ad targeting. “It could be you just had a little fender bender, and you need something to lift you up,” he says. He also says that connected home gadgets, such as thermostats or home automation systems, could allow for creative new ads.

Kiip currently sells ads to brands including General Mills and American Apparel. These ads target coupons on the basis of actions taken in a mobile app, such as completing a game. Jeremy Lockhorn, vice president for emerging media at the ad agency Razorfish, says that kind of promotion could work on the road, too. “There’s potential there for sure,” he says of Kiip’s experiments in the car.

Kiip is far from the only ad company looking for ways to use online-style ad tactics in the real world. Razorfish and some competitors are testing devices that can detect and interact with nearby smartphones, allowing ads to be targeted in retail stores, for example. “That kind of microtargeting becomes the real-world version of the targeting we can do online,” Lockhorn says.

However, Kiip and other companies moving in that direction must find a way to reach people without leaving those people feeling that they no longer control their data. “The user that’s going to interact with your brand really needs to know what they are giving up,” says Jay Giraud, CEO and cofounder of Mojio. “The way I drive my car is personal information.”

NEIL WEBB; COURTESY OF BETH ISRAEL DEACONESS MEDICAL CENTER (HORNG).

MIT TECHNOLOGY REVIEW
VOL. 117 | NO. 4

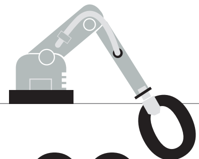

500,000

The number of electric cars that Elon Musk says Tesla's "gigafactory" will be able to produce each year. This year his company sold 35,000 cars in total.

Selling Teslas in China Won't Do Much for the Environment

Because China relies so heavily on coal for power, electric vehicles aren't necessarily an improvement.

By Mike Orcutt

Sales of electric vehicles in China, the world's largest auto market, have been minuscule despite government incentives meant to put five million of the cars on the nation's roads by 2020. Tesla Motors hopes to change that. It has ambitions to break into the Chinese market, and it made its first deliveries of Model S sedans to customers in China in April.

But while having more EVs might help China reach its transportation goals, it probably won't help the environment, given the country's reliance on coal for more than 70 percent of its electricity. Making matters worse, coal in China is often dirtier than it is elsewhere, and many power plants don't employ modern emission-control technologies.

Tesla Model S

Recent research led by Christopher Cherry, a professor of engineering at the University of Tennessee, has shown that in much of China, an electric vehicle the size of a Nissan Leaf accounts for roughly the same amount of carbon dioxide per mile driven as a comparable gasoline-powered car. On top of that, he found that EVs in China account for a larger amount of dangerous particulate emissions than conventional cars.

As in the U.S., electricity production in China is cleaner in some places than in others. Whereas northern China relies almost completely on coal plants, a quarter to a third of the power delivered to central and southern regions comes from hydroelectric dams.

This means that increased sales of EVs are likely to have a mixed environmental impact. Electric vehicles could help clear the air in Chinese cities, many of which have dangerous levels of pollution. Much of that is caused by fine particulate matter emitted from the tailpipes of gas-powered cars. EVs don't have tailpipes, of course.

But this progress could come at a significant environmental cost: higher emissions of particulate matter in rural areas, where most power plants are located. Because burning coal produces far more particulates than burning gasoline, Cherry and his colleagues have found that on average, an EV in China is associated with about 19 times as much fine particulate matter as a comparable gasoline car.

China's government is aiming to clean up its electric grid by increasing its nuclear capacity in the coming years and aggressively deploying renewables. But recent projections suggest that by 2020, coal's share will still be at least as high as 60 percent. The bottom line is that China's grid, and the increasing number of EVs that run off it, will be fairly dirty for a long time.

3 QUESTIONS

Paolo Pirjanian

The CTO of iRobot discusses the company's efforts to build smarter household robots.

You say you want robots to take on more tasks around the home. What technologies must you develop?

The missing link in robotics is low-cost manipulation. We're working on making manipulation much cheaper—for example, using plastic parts, not steel, that can tolerate less precision. Navigation is also a key area, because it allows robots to move around freely and intelligently. We're being helped by the availability of low-cost 3-D sensors. If you combine photos with a 3-D map of a room, you get something like a CAD model. I can say "Mop the kitchen on Tuesdays," or even "Find this book."

Can you really make robots smart enough to do that?

A robot can use the cloud to start learning things about its environment. For example, this object is a cup, and so I have to grab it like this; it looks like it's glass, so I need to grip it tight enough so it doesn't slip but not too hard so it breaks.

What might robots built with this technology do in our homes?

Consumer research tells us that laundry is the number one household task that people spend their time on, so a laundry robot would be on top of the list. But that is a ways off. Through our government and defense business, we have a lot of experience with things that work in rugged outdoor environments, so you can imagine us going into the backyard. —*Tom Simonite*

COURTESY OF TESLA; COURTESY OF IROBOT (PIRJANIAN)

Neuroscience's New

Toolbox

With the invention of optogenetics and other key technologies, researchers can investigate the source of emotions, memory, and consciousness for the first time.

BY
Stephen S. Hall

SCULPTURE
Joshua Harker

What might be called the "make love, not war" branch of behavioral neuroscience began to take shape in (where else?) California several years ago, when researchers in David J. Anderson's laboratory at Caltech decided to tackle the biology of aggression. They initiated the line of research by orchestrating the murine version of Fight Night: they goaded male mice into tangling with rival males and then, with painstaking molecular detective work, zeroed in on a smattering of cells in the hypothalamus that became active when the mice started to fight.

The hypothalamus is a small structure deep in the brain that, among other functions, coördinates sensory inputs—the appearance of a rival, for example—with instinctual behavioral responses. Back in the 1920s, Walter Hess of the University of Zurich (who would win a Nobel in 1949) had shown that if you stuck an electrode into the brain of a cat and electrically stimulated certain regions of the hypothalamus, you could turn a purring feline into a furry blur of aggression. Several interesting hypotheses tried to explain how and why that happened, but there was no way to test them. Like a lot of fundamental questions in brain science, the mystery of aggression didn't go away over the past century—it just hit the usual empirical roadblocks. We had good questions but no technology to get at the answers.

By 2010, Anderson's Caltech lab had begun to tease apart the underlying mechanisms and neural circuitry of aggression in their pugnacious mice. Armed with a series of new technologies that allowed them to focus on individual clumps of cells within brain regions, they stumbled onto a surprising anatomical discovery: the tiny part of the hypothalamus that seemed correlated with aggressive behavior was intertwined with the part associated with the impulse to mate. That small duchy of cells—the technical name is the ventromedial hypothalamus—turned out to be an assembly of roughly 5,000 neurons, all marbled together, some of them seemingly connected to copulating and others to fighting.

"There's no such thing as a generic neuron," says Anderson, who estimates that there may be up to 10,000 distinct classes of neurons in the brain. Even tiny regions of the brain contain a mixture, he says, and these neurons "often influence behavior in different, opposing directions." In the case of the hypothalamus, some of the neurons seemed to become active during aggressive behavior, some of them during mating behavior, and a small subset—about 20 percent—during both fighting and mating.

That was a provocative discovery, but it was also a relic of old-style neuroscience. Being active was not the same as causing the behavior; it was just a correlation. How did the scientists know for sure what was triggering the behavior? Could they provoke a mouse to pick a fight simply by tickling a few cells in the hypothalamus?

A decade ago, that would have been technologically impossible. But in the last 10 years, neuroscience has been transformed by a remarkable new technology called optogenetics, invented by scientists at Stanford University and first described in 2005. The Caltech researchers were able to insert a genetically modified light-sensitive gene into specific cells at particular locations in the brain of a living, breathing, feisty, and occasionally canoodling male mouse. Using a hair-thin fiber-optic thread inserted into that living brain, they could then turn the neurons in the hypothalamus on and off with a burst of light.

Anderson and his colleagues used optogenetics to produce a video dramatizing the love-hate tensions deep within rodents. It shows a male mouse doing what comes naturally, mating with a female, until the Caltech researchers switch on the light, at which instant the murine lothario flies into a rage. When the light is on, even a mild-mannered male mouse can be induced to attack whatever target happens to be nearby—his reproductive partner, another male mouse, a castrated male (normally not perceived as a threat), or, most improbably, a rubber glove dropped into the cage.

"Activating these neurons with optogenetic techniques is sufficient to activate aggressive behavior not only toward appropriate targets like another male mouse but also toward inappropriate targets, like females and even inanimate objects," Anderson says. Conversely, researchers can inhibit these neurons in the middle of a fight by turning the light off, he says: "You can stop the fight dead in its tracks."

Moreover, the research suggests that lovemaking overrides war-making in the calculus of behavior: the closer a mouse was to consummation of the reproductive act, the more resistant (or oblivious) he became to the light pulses that normally triggered aggression. In a paper published in *Biological Psychiatry*, titled "Optogenetics, Sex, and Violence in the Brain: Implications for Psychiatry," Anderson noted, "Perhaps the imperative to 'make love, not war' is hard-wired into our nervous system, to a greater extent than we have realized." We may be both lovers and fighters, with the slimmest of

PREVIOUS SPREAD: PHOTOGRAPH BY BRUCE PETERSON

Optogenetics:
Light Switches for Neurons

KEY PLAYERS

Gero Miesenböck
University of Oxford

Edward Boyden
MIT

Karl Deisseroth
Stanford University

Feng Zhang
Broad Institute of MIT and Harvard

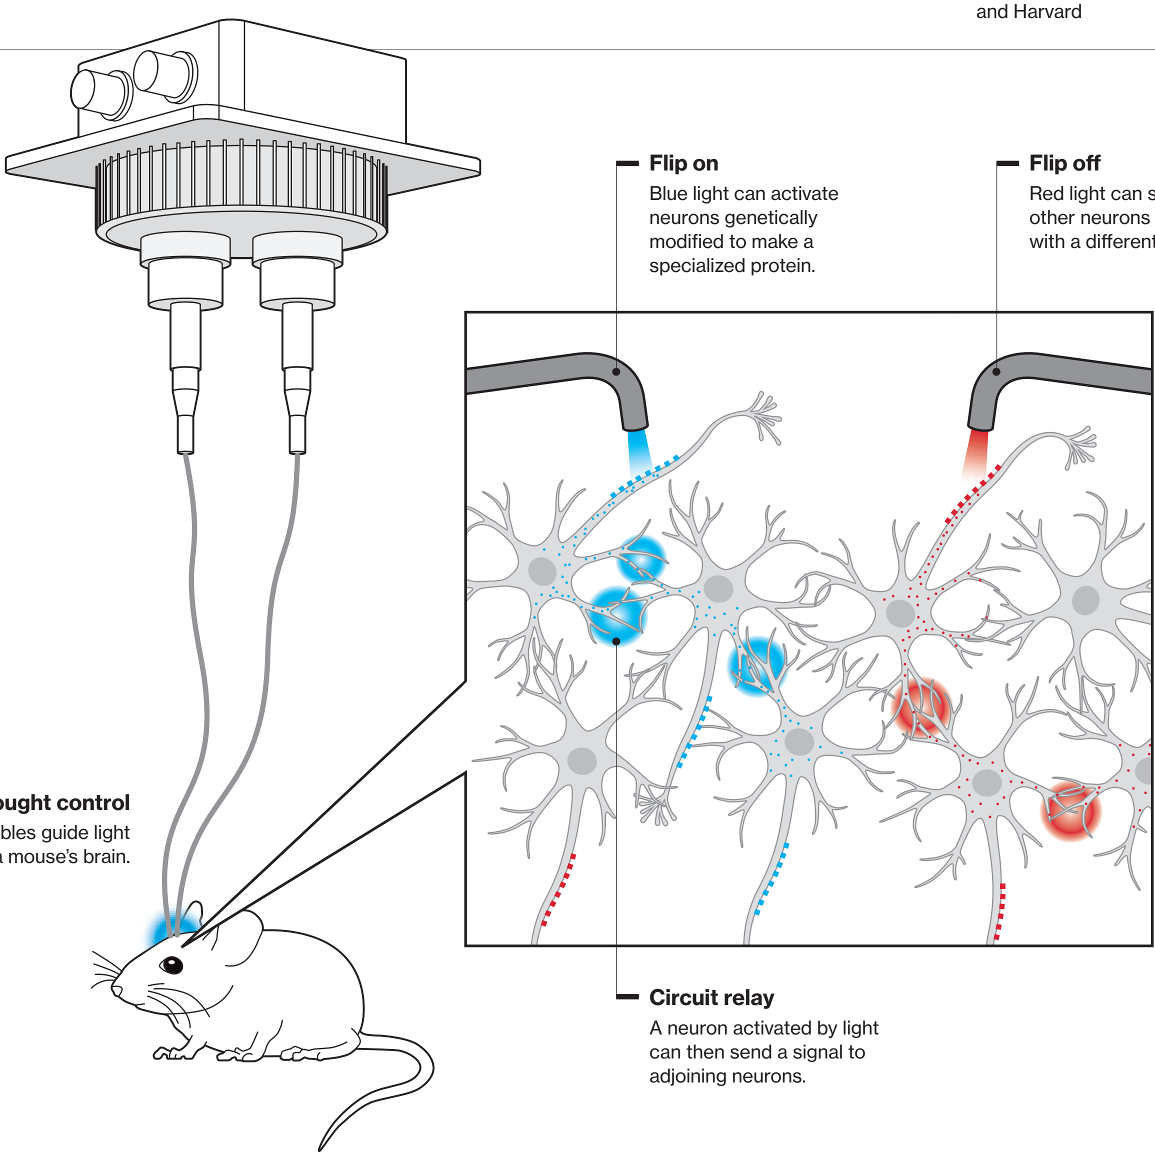

Flip on
Blue light can activate neurons genetically modified to make a specialized protein.

Flip off
Red light can silence other neurons modified with a different protein.

Thought control
Thin cables guide light into a mouse's brain.

Circuit relay
A neuron activated by light can then send a signal to adjoining neurons.

Among the billions of neurons in a human brain, what role do specific neurons play in controlling movement or cognition, or causing disorders like depression and autism? Figuring out the functions of various neurons and circuits in mice or other lab animals can offer answers.

A technique called optogenetics gives scientists the power to turn neurons on or off. Scientists can then explore whether a particular set of neurons is responsible for a behavior or disorder. With optogenetics, researchers can manipulate neurons in worms, flies, mice, and even monkeys. The technology has been used to study neural processes underlying epilepsy, addiction, depression, and more.

Optogenetics requires genetic modifications to neurons so that they produce light-reactive proteins, as well as a source of wavelength-tuned light that is often delivered through a hole in the skull. For some time, the light-sensitive proteins were most efficient at activating neurons; silencing occurred only slowly and weakly. But recently, researchers engineered proteins that can also efficiently silence neurons, expanding the toolkit for studying the roles that different neurons play in the brain's many circuits.

While it's unlikely to be used in human brains anytime soon, optogenetics can in the short term point the way to new treatment strategies for devastating brain disorders.

GRAPHIC BY JOHN MACNEILL

Growing Neurons:
Studying What Goes Wrong

1 in 4
Proportion of people worldwide who will have a mental disorder in their lifetime.

5
Number of major brain disorders, including schizophrenia, that share genetic variations.

400
Number of genes or chromosomal regions that may be involved in autism risk.

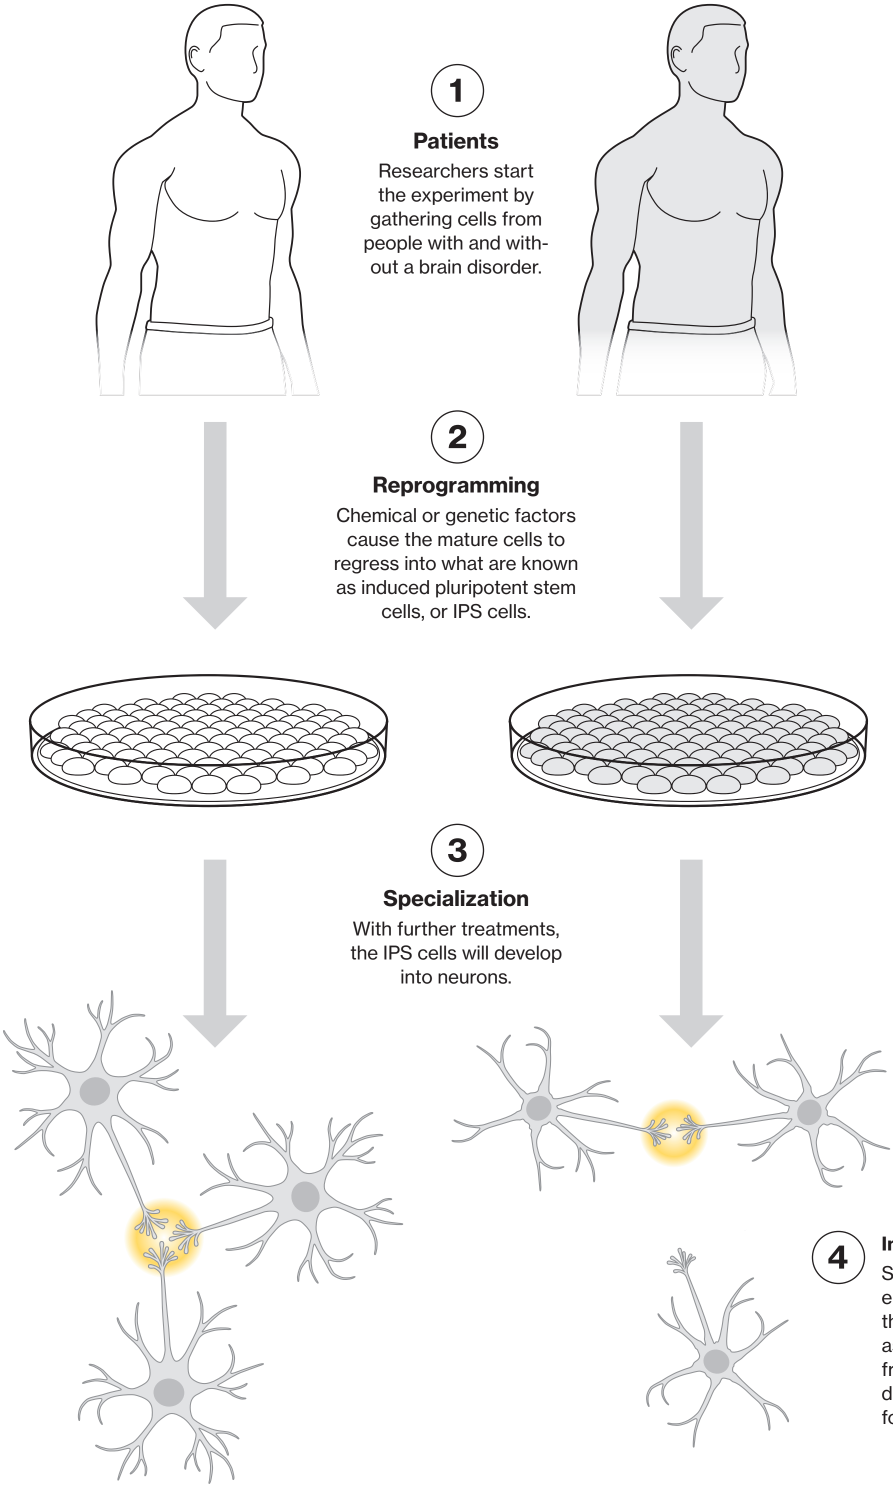

Inspection
Scientists look for differences between neurons in the two populations, such as whether those derived from the person with the disorder are smaller or form fewer connections.

The cellular and genetic changes that drive most brain disorders aren't well known. But new technologies are making it possible to understand these changes, which could help scientists identify new drugs or other treatments.

For example, scientists can grow a patient's neurons in a petri dish. All the scientists need to begin are skin cells or another easily accessed type of cell; these cells are transformed first into stem cells and then into neurons. The neurons in the dish can be examined for changes in appearance and function. Scientists can also test drugs on these neurons to see whether any compounds ameliorate the defects.

Scientists can also use genome editing to introduce into stem cells the precise changes that are believed to be involved in a disorder. They can then check whether the resulting neurons are abnormal.

GRAPHIC BY JOHN MACNEILL

neurological distances separating the two impulses.

No one is suggesting that we're on the verge of deploying neural circuit breakers to curb aggressive behavior. But, as Anderson points out, the research highlights a larger point about how a new technology can reinvent the way brain science is done. "The ability of optogenetics to turn a largely correlational field of science into one that tests causation has been transformative," he says.

What's radical about the technique is that it allows scientists to perturb a cell or a network of cells with exquisite precision, the key to sketching out the circuitry that affects various types of behavior. Whereas older technologies like imaging allowed researchers to watch the brain in action, optogenetics enables them to influence that action, tinkering with specific parts of the brain at specific times to see what happens.

Optogenetics and other new techniques mean scientists can begin to pinpoint the function of the thousands of different types of neurons among the roughly 86 billion in the human brain.

And optogenetics is just one of a suite of revolutionary new tools that are likely to play leading roles in what looks like a heyday for neuroscience. Major initiatives in both the United States and Europe aspire to understand how the human brain—that tangled three-pound curd of neurons, connective tissue, and circuits—gives rise to everything from abstract thought to basic sensory processing to emotions like aggression. Consciousness, free will, memory, learning—they are all on the table now, as researchers use these tools to investigate how the brain achieves its seemingly mysterious effects (see "Searching for the "Free Will" Neuron," page 64).

Connections

More than 2,000 years ago, Hippocrates noted that if you want to understand the mind, you must begin by studying the brain. Nothing has happened in the last two millennia to change that imperative—except the tools that neuroscience is bringing to the task.

The history of neuroscience, like the history of science itself, is often a story of new devices and new technologies. Luigi Galvani's first accidental electrode, which provoked the twitch of a frog's muscle, has inspired every subsequent electrical probe, from Walter Hess's cat prod to the current therapeutic use of deep brain stimulation to treat Parkinson's disease (approximately 30,000 people worldwide now have electrodes implanted in their brains to treat this condition). The patch clamp allowed neuroanatomists to see the ebb and flow of ions in a neuron as it prepares to fire. And little did Paul Lauterbur realize, when he focused a strong magnetic field on a single hapless clam in his lab at the State University of New York at Stony Brook in the early 1970s, that he and his colleagues were laying the groundwork for the magnetic resonance imaging (MRI) machines that have helped reveal the internal landscape and activity of a living brain.

But it is the advances in genetics and genomic tools during the last few years that have truly revolutionized neuroscience. Those advances made the genetic manipulations at the heart of optogenetics possible. Even more recent genome-editing methods can be used to precisely alter the genetics of living cells in the lab. Along with optogenetics, these tools mean scientists can begin to pinpoint the function of the thousands of different types of nerve cells among the roughly 86 billion in the human brain.

Nothing testifies to the value of a new technology more than the number of scientists who rapidly adopt it and use it to claim new scientific territories. As Edward Boyden, a scientist at MIT who helped develop optogenetics, puts it, "Often when a new technology comes out, there's a bit of a land grab."

And even as researchers grab those opportunities in genomics and optogenetics, still other advances are coming on the scene. A new chemical treatment is making it possible to directly see nerve fibers in mammalian brains; robotic microelectrodes can eavesdrop on (and perturb) single cells in living animals; and more sophisticated imaging techniques let researchers match up nerve cells and fibers in brain slices to create a three-dimensional map of the connections. Using these tools together to build up an understanding of the brain's activity, scientists hope to capture the biggest of cognitive game: memory, decision-making, consciousness, psychiatric illnesses like anxiety and depression, and, yes, sex and violence.

In January 2013, the European Commission invested a billion euros in the launch of its Human Brain Project, a 10-year initiative to map out all the connections in the brain. Several months later, in April 2013, the Obama administration announced an initiative called Brain Research through Advanced Innovative Neurotechnologies (BRAIN), which is expected to pour as much as $1 billion into the field, with much of the early funding earmarked for technology development. Then there is the Human Connectome Project, which aims to use electron microscope images of sequential slices of brain tissue to map nerve cells and their connections in three dimensions. Complementary connectome and mapping initiatives are getting under way at the Howard Hughes Medical Institute in Virginia and the Allen Institute for Brain Science in Seattle. They are all part of a large global effort, both publicly and privately funded, to build a comprehensive picture of the human brain, from the level of genes and cells to that of connections and circuits.

Last December, as an initial step in the BRAIN Initiative, the National Institutes of Health solicited proposals for $40 million worth of projects on technology development in the neurosciences. "Why is the BRAIN Initiative putting such a heavy emphasis on technology?" says Cornelia Bargmann, the Rockefeller University neuroscientist who co-directs the planning process for the project. "The real goal is to understand how the brain works, at many levels, in space and time, in many different neurons, all at once. And what's prevented us from understanding that is limitations in technology."

Scientists speculated that if you could smuggle the gene for a light-sensitive protein into a neuron and then pulse the cell with light, you might trigger it to fire. You could turn specific neurons on and off.

Eavesdropping

Optogenetics had its origins in 2000, in late-night chitchat at Stanford University. There, neuroscientists Karl Deisseroth and Edward Boyden began to bounce ideas back and forth about ways to identify, and ultimately manipulate, the activity of specific brain circuits. Deisseroth, who had a PhD in neuroscience from Stanford, longed to understand (and someday treat) the mental afflictions that have vexed humankind since the era of Hippocrates, notably anxiety and depression (see "Shining Light on Madness," page 34). Boyden, who was pursuing graduate work in brain function, had an omnivorous curiosity about neurotechnology. At first they dreamed about deploying tiny magnetic beads as a way to manipulate brain function in intact, living animals. But at some point during the next five years, a different kind of light bulb went off.

Since the 1970s, microbiologists had been studying a class of light-sensitive molecules known as rhodopsins, which had been identified in simple organisms like bacteria, fungi, and algae. These proteins act like gatekeepers along the cell wall; when they detect a particular wavelength of light, they either let ions into a cell or, conversely, let ions out of it. This ebb and flow of ions mirrors the process by which a neuron fires: the electrical charge within the nerve cell builds up until the cell unleashes a spike of electrical activity flowing along the length of its fiber (or axon) to the synapses, where the message is passed on to the next cell in the pathway. Scientists speculated that if you could smuggle the gene for one of these light-sensitive proteins into a neuron and then pulse the cell with light, you might trigger it to fire. Simply put, you could turn specific neurons in a conscious animal on—or off—with a burst of light.

In 2004, Deisseroth successfully inserted the gene for a light-sensitive molecule from algae into mammalian neurons in a dish. Deisseroth and Boyden went on to show that blue light could induce the neurons to fire. At about the same time, a graduate student named Feng Zhang joined Deisseroth's lab. Zhang, who had acquired a precocious expertise in the techniques of both molecular biology and gene therapy as a high school student in Des Moines, Iowa, showed that the gene for the desired protein could be introduced into neurons by means of genetically engineered viruses. Again using pulses of blue light, the Stanford team then demonstrated that it could turn electrical pulses on and off in the virus-modified mammalian nerve cells. In a landmark paper that appeared in *Nature Neuroscience* in 2005 (after, Boyden says, it was rejected by *Science*), Deisseroth, Zhang, and Boyden described the technique. (No one would call it "optogenetics" for another year.)

Neuroscientists immediately seized on the power of the technique by inserting light-sensitive genes into living animals. Researchers in Deisseroth's own lab used it to identify new pathways that control anxiety in mice, and both Deisseroth's team and his collaborators at Mount Sinai Hospital in New York used it to turn depression on and off in rats and mice. And Susumu Tonegawa's lab at MIT recently used optogenetics to create false memories in laboratory animals.

When I visited Boyden's office at MIT's Media Lab last December, the scientist called up his favorite recent papers involving optogenetics. In a rush of words as rapid as his keystrokes, Boyden described second-generation technologies already being developed. One involves eavesdropping on single nerve cells in anesthetized and conscious animals in order to see "the things roiling underneath the sea of activity" within a neuron when the animal is unconscious. Boyden said, "It literally sheds light on what it means to have thoughts and awareness and feelings."

Boyden's group had also just sent off a paper reporting a new twist on optogenetics: separate, independent neural pathways can be perturbed simultaneously with red and blue wavelengths of

Brain Mapping: Charting the Information Superhighways

86
BILLION
Number of neurons in the adult brain.

10,000
Upper range of the number of synapses, or connections, one neuron can form.

1,000
TRILLION
An estimate of the number of synapses in the human brain.

Many of the brain's functions result from processes in which bundles of neurons work together to transmit signals between different regions. Mapping these pathways will help scientists understand how the brain works. The task is immense, but many teams are contributing to the effort, looking at different aspects of the brain's anatomy.

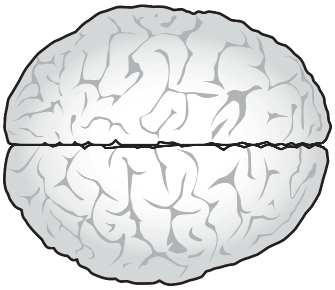

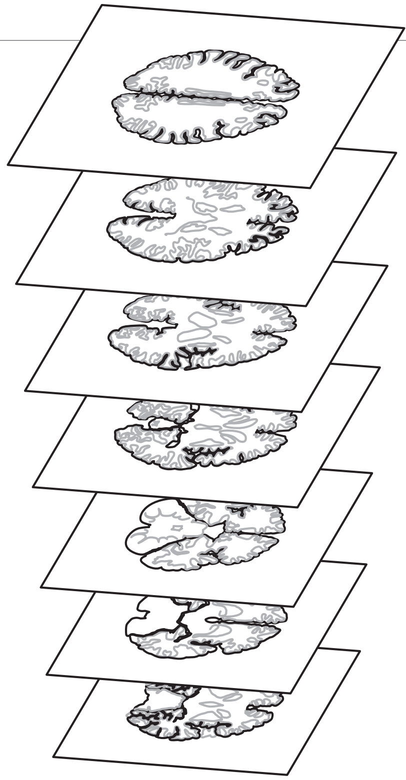

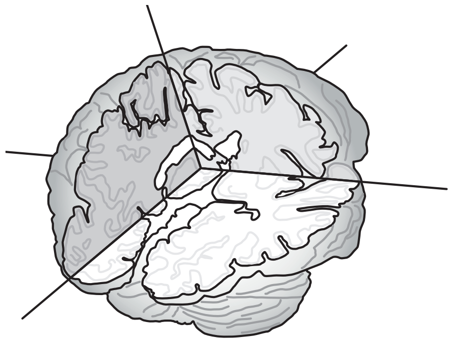

1

Brain prep

At the Jülich Research Centre in Germany, a human brain is removed during autopsy, chemically preserved, and then scanned in an MRI machine to help guide virtual reconstruction.

2

Slice and scan

The brain is sliced into ultrathin slivers that can be stained and scanned.

3

Stack and map

The individual images of slivers are computationally processed into a three-dimensional model.

Other Mapping Technologies

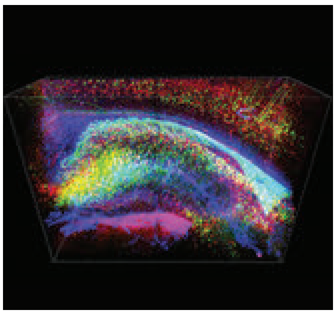

Clarity

By removing the fats from mammalian brains, scientists can stain and image fine details of neurons, as seen in this mouse brain.

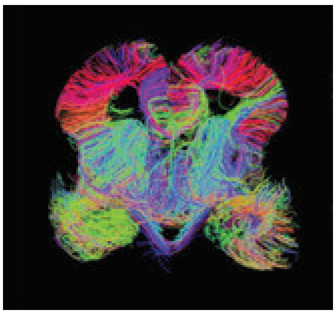

BrainSpan Atlas

The Allen Institute for Brain Science has created three-dimensional maps of structures and gene activity in prenatal brains.

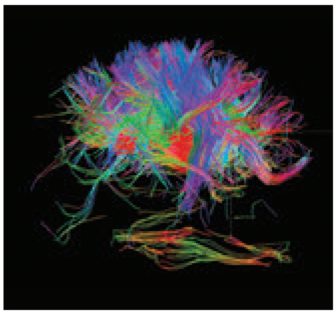

The NIH Human Connectome Project

An advanced form of MRI enables scientists to visualize microscopic cell structures and large brain structures simultaneously.

GRAPHIC BY JOHN MACNEILL; COURTESY OF THE DEISSEROTH LAB (CLARITY); COURTESY OF ALLEN INSTITUTE FOR BRAIN SCIENCE (BRAINSPAN); COURTESY OF HUMAN CONNECTOME PROJECT light. The technique has the potential to show how different circuits interact with and influence each other. His group is also working on "insanely dense" recording probes and microscopes that aspire to capture whole-brain activity. The ambitions are not modest. "Can you record all the cells in the brain," he says, "so that you can watch thoughts or decisions or other complex phenomena emerge as you go from sensation to emotion to decision to action site?"

A few blocks away, Feng Zhang, who is now an assistant professor at MIT and a faculty member at the Broad Institute, listed age-old neuroscience questions that might now be attacked with the new technologies. "Can you do a memory upgrade and increase the capacity?" he asked. "How are neural circuits genetically encoded? How can you reprogram the genetic instructions? How do you fix the genetic mutations that cause miswiring or other malfunctions of the neural system? How do you make the old brain younger?"

Scientists often employ words like "surprising" and "unexpected" to characterize recent results, reflecting the impact optogenetics has had on the understanding of psychiatric illnesses.

In addition to helping to invent optogenetics, Zhang played a central role in developing a gene-editing technique called CRISPR (see "10 Breakthrough Technologies: Genome Editing," May/June). The technology allows scientists to target a gene—in neurons, for example—and either delete or modify it. If it's modified to include a mutation known or suspected to cause brain disorders, scientists can study the progression of those disorders in lab animals. Alternatively, researchers can use CRISPR in the lab to alter stem cells that can then be grown into neurons to see the effects.

Transparency

Back at Stanford, when he's not seeing patients with autism spectrum disorders or depression in the clinic, Deisseroth continues to invent tools that he and others can use to study these conditions. Last summer, his lab reported a new way for scientists to visualize the cables of nerve fibers, known as "white matter," that connect distant precincts of the brain. The technique, called Clarity, first immobilizes biomolecules such as protein and DNA in a plastic-like mesh that retains the physical integrity of a postmortem brain. Then researchers flush a kind of detergent through the mesh to dissolve all the fats in brain tissue that normally block light. The brain is rendered transparent, suddenly exposing the entire three-dimensional wiring pattern to view.

Together, the new tools are transforming many conventional views in neuroscience. For example, as Deisseroth noted in a review article published earlier this year in *Nature*, optogenetics has challenged some of the ideas underlying deep brain stimulation, which has been widely used to treat everything from tremors and epilepsy to anxiety and obsessive-compulsive disorder. No one knows just why it works, but the operating assumption has been that its therapeutic effects derive from electrical stimulation of very specific brain regions; neurosurgeons exert extraordinary effort to place electrodes with the utmost precision.

In 2009, however, Deisseroth and colleagues showed that specifically stimulating the white matter, the neural cables that happen to lie near the electrodes, produced the most robust clinical improvement in symptoms of Parkinson's disease. In other words, it wasn't the neighborhood of the brain that mattered so much as which neural highways happened to pass nearby. Scientists often employ words like "surprising" and "unexpected" to characterize such recent results, reflecting the impact that optogenetics has had on the understanding of psychiatric illness.

In the same vein, Caltech's Anderson points out that the public and scientific infatuation with functional MRI studies over the last two decades has created the impression that certain regions of the brain act as "centers" of neural activity—that the amygdala is the "center" of fear, for example, or the hypothalamus is the "center" of aggression. But he likens fMRI to looking down on a nighttime landscape from an airplane at 30,000 feet and "trying to figure out what is going on in a single town." Optogenetics, by contrast, has provided a much more detailed view of that tiny subdivision of cells in the hypothalamus, and thus a much more complex and nuanced picture of aggression. Activating specific neurons in that little town can tip an organism to make war, but activating the neurons next door can nudge it to make love.

The new techniques will give scientists the first glimpses of human cognition in action—a look at how thoughts, feelings, forebodings, and dysfunctional mental activity arise from the neural circuitry and from the activity of particular types of cells. Researchers are just beginning to gain these insights, but given the recent pace of technology development, the picture might emerge sooner than anyone dreamed possible when the light of optogenetics first flickered on a few years ago. ▪

Stephen S. Hall is a science writer and author in New York City. His last feature for MIT Technology Review *was "Repairing Bad Memories" in July/August 2013.*

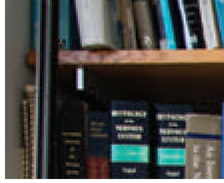

Q+A
Joseph LeDoux

When it comes to the study of memory, we might be living in something of a golden age. Researchers are exploring provocative questions about what memory fundamentally is—and how it might be manipulated. Some scientists are tweaking the brains of lab rats in order to implant false memories or remove specific memories. Others are looking into how memory might be enhanced. Such research often sounds creepy, but it could lead to ways of staving off dementia, neutralizing post-traumatic stress disorder, reducing anxiety, treating depression, or curbing addiction.

Much of this work is possible because neuroscientists have realized that memory is more plastic than previously thought. Think of something that you did long ago—on a sunny afternoon when you were a child, let's say. Does your brain rummage around for that memory, show it to you, and put it back intact, as you might do with a photograph in a trunk in the attic? For decades, the prevailing answer was essentially yes—that strong memories were "consolidated" in the brain and remained static. But it now appears the opposite is true: every time you remember something, your brain rewrites or "reconsolidates" the memory. Your memory of any sunny day in your childhood is merely a version of the last time you thought about it.

Among the stunning implications is that intervening in the reconsolidation process can alter a memory and change how it feels. Some of the most intriguing research on this idea has been led by Joseph LeDoux, a neuroscientist who has been working since the 1970s to investigate how processes in the brain generate emotions. In recent years, he and colleagues have investigated whether giving people an antianxiety drug as they recall a traumatic experience can reduce the dread they feel upon further recollections. If it works, it could be one of many opportunities for reshaping memory, as LeDoux told *MIT Technology Review*'s deputy editor, Brian Bergstein, in his NYU office.

BB: Giving people a pill as they reconsolidate a troubling memory sounds like an elegant treatment for PTSD and other disorders. How well has the research progressed?

JL: The idea is still viable. The drug that's going to do the trick in humans—we don't yet know what it is.

BB: Why not the antianxiety drug propranolol, which has been used in several experiments?

JL: We know this thing works very well in rats. [In humans] there's been some success, but it doesn't seem that's actually going to be a very robust solution.

BB: What are the prospects for improving the memory of people with dementia—or any of us, for that matter?

JL: The issue you get at whenever you're dealing with a memory failure of any kind is: is the problem an inability to retrieve a memory that's there, or is it that the memory is no longer there and by no means could you retrieve it? And it's definitely the case that we

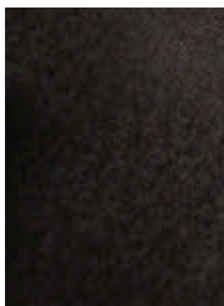

LeDoux researches how mechanisms in the brain work to make us who we are.

Adam Lerner

have retrieval problems, right? You can't remember something, you kind of know it's there, and then a couple of hours later it pops up. There is stuff in our brain that we can't easily retrieve. I think that's where there's hope. You can do things that would facilitate the retrieval process.

BB: What might that be?

JL: The simplest idea would be that it's a problem of low arousal in the brain. We know that emotionally arousing situations are more likely to be remembered than mundane ones. A big part of the reason for this is that in significant situations chemicals called neuromodulators are released, and they enhance the memory storage process. The brain is more alert and attentive. All the spokes are working and all the gears are oiled.

BB: So could we develop memory vitamins that offer that same boost?

JL: [The effects of the neuromodulators] can be mimicked with drugs that do the same thing. Or if you want to remember something, put the information in a significant context. In other words, make it meaningful by thinking of it in a way that adds some positive or negative charge, or by doing something that increases your level of arousal—exercise, for example.

BB: Would a memory prosthetic be possible—something put into the brain to restore lost memories in someone with dementia or a brain injury?

JL: DARPA [the U.S. military's R&D agency] seems to be going full steam ahead on these kinds of technologies. What they plan to do is put chips in [the brain]. It would be like a prosthesis—instead of moving your arm, you're fixing memory. I have no idea how they would achieve that.

BB: We don't know the route to get there?

JL: *I* don't know the route.

Memory is, first of all, not in one spot—it's distributed across probably multiple brain areas, many millions of synapses in those brain areas. So I don't know how you reconfigure it, fix it. I don't know how you regenerate the right patterns of connectivity. I don't think you could restore lost memories. But what you might conceivably do is restore some ability to store new memories.

"You can't remember something, you kind of know it's there, and then a couple of hours later it pops up. There is stuff in our brain that we can't easily retrieve. I think that's where there's hope. You can do things that would facilitate the retrieval process."

BB: Messing with memory is a huge deal. It goes to the core of who we are. Treating PTSD would be wonderful, but isn't it also possible to make people artificial Pollyannas?

JL: Or fearless monsters. There's always going to be ethical implications. But we'll just have to sort that out.

When we first published this work [on reconsolidation], someone wrote a commentary in the *New York Times* saying, "Let's say you were a Holocaust survivor. You lived 50 years with these awful memories, and all of a sudden you're erasing memories of the Holocaust. What would that do to your personality? It's who you are now." [After further research], the conclusion that we came up with is that a patient and therapist would have to slowly chip away at a memory to a level they were comfortable with. [And] the research so far suggests it reduces the zing, takes the emotional valence out of the situation, rather than erases the memory itself.

The other side of that is you can also intensify memories. So we did studies in rats where we give them propranolol, and that weakens the memory. But if you give them isoproterenol [which has the opposite effect on the brain's neurostimulators], the memory is now stronger. If you have some way to strengthen memory after retrieval, it could be stronger and better.

BB: What's a situation in which people would want to do that?

JL: In general, people who have sluggish memories, they're not forming memories very well. Were they to be on a low dose of isoproterenol, they might get a little extra boost.

The other idea is that we might be able to give people positive experiences and a shot of isoproterenol, and store those instead of negative experiences.

BB: You could build up more positive memories in them?

JL: Right, right.

BB: I imagine that could be useful in helping people struggling with depression.

JL: And as far as I know it's never been done. It's doable.

BB: So why hasn't it been done?

JL: Well, you know, I had the idea a while ago, and then I kind of forgot about it. And now the memory of it—maybe I'll think about doing something. *[He laughs.]* T

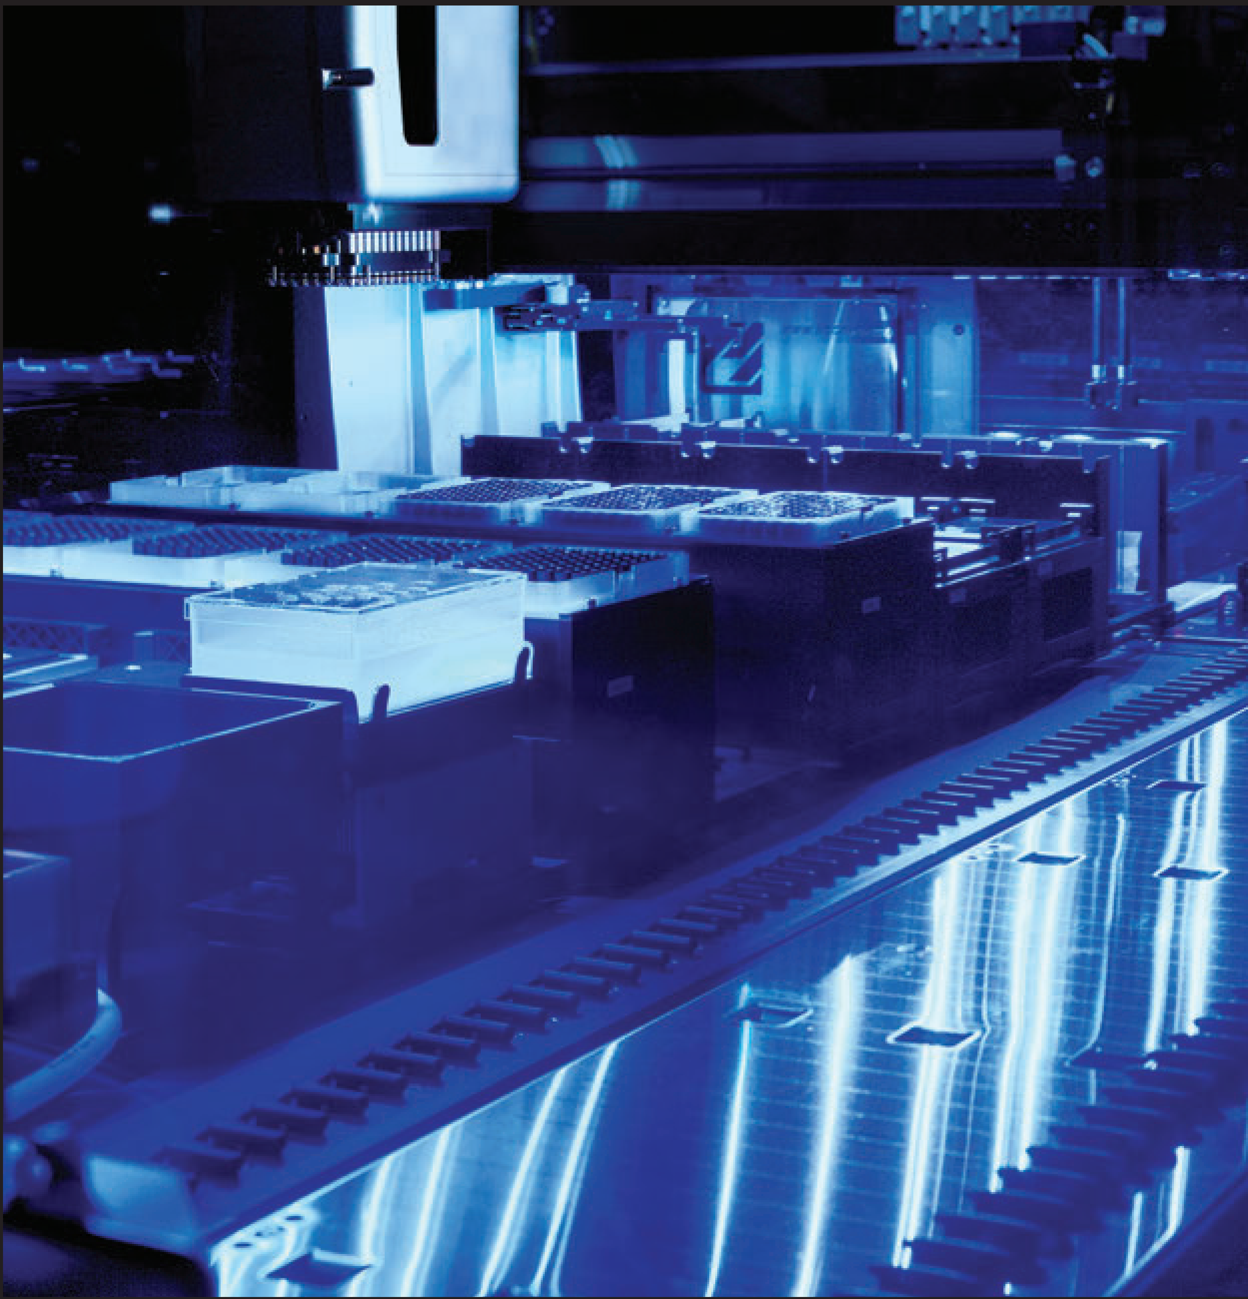

Shining Light on Madness

It's been decades since there was a true breakthrough in drugs for psychiatric illness.

But better research tools and new insights into the genetics of brain disorders could revive the moribund effort to improve treatment.

By David Rotman

AT NOVARTIS'S RESEARCH LAB IN CAMbridge, Massachusetts, a large incubator-like piece of equipment is helping give birth to a new era of psychiatric drug discovery. Inside it, bathed in soft light, lab plates hold living human stem cells; robotic arms systematically squirt nurturing compounds into the plates. Thanks to a series of techniques perfected over the last few years in labs around the world, such stem cells—capable of developing into specialized cell types—can now be created from skin cells. When stem cells derived from people with, say, autism or schizophrenia are grown inside the incubator, Novartis researchers can nudge them to develop into functioning brain cells by precisely varying the chemicals in the cell cultures.

They're not exactly creating schizophrenic or autistic neurons, because the cells aren't working within the circuitry of the brain, but for drug-discovery purposes it's the next best thing. For the first time, researchers have a way to directly examine in molecular detail what's going wrong in the brain cells of patients with these illnesses. And, critically for the pharmaceutical company, there is now a reliable method of screening for drugs that might help. Do the neurons look different from normal ones? Is there a flaw in the way they form connections? Could drugs possibly correct the abnormalities? The answer to each of these questions is a very preliminary yes.

The technique is so promising that Novartis has resumed trying to discover new psychiatric drugs after essentially abandoning the quest. What's more, it's been introduced at a time when knowl-

An incubator at Novartis's labs is used to grow stem cells. Scientists can derive such cells from patients and, under the right conditions, prompt them to develop into neurons for research and drug screening.

Photographs by Adam DeTour

edge about the genetics behind brain disorders is expanding rapidly and other new tools, including optogenetics and more precise genome editing (see "Neuroscience's New Toolbox," page 20), are enabling neuroscientists to probe the brain directly. All these developments offer renewed hope that science could finally deliver more effective treatments for the millions of people beset by devastating brain disorders.

A revival in psychiatric drug development is badly needed: there hasn't been a breakthrough medicine for any of the common mental illnesses, including schizophrenia, bipolar disorder, or severe depression, in roughly 50 years. From the late 1940s through the 1960s, a series of serendipitous discoveries, beginning with the finding that lithium could help bipolar patients, transformed the treatment of the mentally ill. It became possible to quiet the hallucinations and delusions of schizophrenia and offer a drug to the severely depressed. The sudden availability of pharmacological relief transformed psychiatry and played a role in closing down many of the mammoth mental hospitals of the era. But then, almost as suddenly as it had started, the revolution stalled.

Many of the drugs discovered in the 1950s and 1960s are still the most effective treatments available for schizophrenia, anxiety disorders, and depression. But while these medications have improved the lives of some patients, they are ineffective for others, and they are woefully inadequate in treating many of the worst symptoms. What's more, the drugs can have severe side effects.

Take schizophrenia, for example. Existing antipsychotic drugs can make the hallucinations and delusions disappear, but they don't improve the so-called negative symptoms—the disruption of emotions such as pleasure, which can leave people uninterested in communicating or even in living. Existing drugs also have no effect on the way schizophrenia can impair concentration, decision-making, and working memory (critical in such tasks as language comprehension). These debilitating cognitive problems make it impossible for people to work and difficult for them even to make the seemingly simple logical choices involved in everyday life. Insidiously, such symptoms can strike high-performing individuals, often in their late teens. "People don't understand," says Guoping Feng, a professor of neuroscience at MIT who studies the neural basis of psychiatric disorders. "They ask, once a patient is given antipsychotic medicine, 'Why can't you go to work?' But [those with schizophrenia] can't work because they don't have cognitive functions, they don't have normal executive functions. And there are no drugs for this." On top of that are the side effects of antipsychotic medicines, which can include Parkinson's-like movement disorders, dramatic weight gain, or a potentially deadly loss of white blood cells. In short, the illness destroys many patients' lives.

Finally, many people with brain disorders are simply not helped at all by available drugs. Antidepressants work well for some people but do nothing for many others, and there are no effective drug treatments for the social disabilities or repetitive behaviors caused by autism.

Overall, neuropsychiatric illness is a leading cause of disability. According to the National Institute of Mental Health (NIMH) in Rockville, Maryland, 26 percent of American adults suffer from a "diagnosable mental disorder" in any given year. Severe depression, the most common of these disorders, is the leading cause of disability in the U.S. for individuals between 15 and 44. Around 1 percent of the American population suffers from schizophrenia; one in 68 American children is diagnosed with an autism spectrum disorder.

Though the need for better treatments is unquestionable, drug companies had until very recently simply run out of good ideas. The drugs developed in the 1950s and 1960s were discovered by accident, and no one knew how or why they worked. In the subsequent decades, drug researchers reverse-engineered the medi-

The Worldwide Cost

Disability-adjusted life years (DALY) is a commonly used public-health metric meant to express the number of years lost as a result of bad health, disability, or premature death.

Disorder	DALYs
Unipolar depressive disorders	65,000,000
Bipolar affective disorder	14,000,000
Schizophrenia	17,000,000
Epilepsy	8,000,000
Alcohol use disorders	24,000,000
Alzheimer's and other dementias	11,000,000
Drug use disorders	8,000,000
Post-traumatic stress disorder	3,000,000
Obsessive-compulsive disorder	5,000,000

Sources: Harvard School of Public Health and the World Economic Forum (2011)

"We were led down a path that said depression is about being a quart low in serotonin, and schizophrenia means you have a bit too much dopamine on board. But that just isn't how the brain works. The brain isn't a bowl of soup."

cations to identify the brain molecules that the drugs acted on, such as dopamine and serotonin. In retrospect, however, scientists now realize that while tweaking the levels of these chemicals addressed some symptoms of psychiatric disorders, it was a crude strategy that ignored the biological mechanisms underlying the illnesses.

"By studying the drugs, we were led down a path that said depression is about being a quart low in serotonin, and schizophrenia means you have a bit too much dopamine on board," says Thomas Insel, the director of NIMH. "But that just isn't how the brain works. The brain isn't a bowl of soup; it's really a complex network of networks." Psychiatric illnesses such as schizophrenia, severe depression, and bipolar disorder, Insel says, "are fundamentally disorders of the brain circuits." It is only in the last few years, he adds, that technologies such as optogenetics have allowed neuroscientists to shift their thinking from "soups to sparks"—electrical impulses—and begin exploring the circuits involved in brain disorders. Still, he says, the search for treatments is slowed by "our profound ignorance of the brain."

Sales of Psychiatric Drugs

Despite a dearth of effective medicines for many brain disorders, demand is very strong.

Drug class or indication	Global sales
Antipsychotics	$22 billion
Antidepressants	$20 billion
Antianxiety	$11 billion
Stimulants	$5.5 billion
Dementia	$5.5 billion
Sleep disorders	$4.5 billion
Substance use and addictive disorders	$3 billion

Source: Science Translational Medicine (2012), citing 2010 sales

Another obvious impediment to finding better drugs is that there haven't been reliable ways to screen them. Because researchers have had limited ability to measure how potential psychiatric drugs affect the biology of lab animals, they've "concocted" tests based on the way existing drugs affect animal behavior, says Steven Hyman, director of the Stanley Center for Psychiatric Research at the Broad Institute of Harvard and MIT. One conventional assay for antidepressants, for example, is called the "forced swim test." When rats given the commonly used drug imipramine, which was invented in the 1950s and is still considered one of the most effective medicines for depression, are dropped in a bucket of cold water, they swim longer before giving up. The animals' propensity to stop struggling has been rationalized as a measure of "behavioral despair," Hyman says, but there's actually no proof that the behavior in the test reflects human depression. Though the swim test has been used for 50 years to test antidepressants and is still widely used, all it probably does is select for drugs that mimic the effects of imipramine in allowing a rodent to swim longer, he says. That has led to a series of "me-too drugs."

The discovery of new psychiatric drugs is "dangerously stalled," Hyman says: in terms of efficacy, antidepressants "maxed out" in the 1950s and antipsychotics in the 1960s. Though a number of new psychiatric drugs have been marketed in recent decades, says Richard A. Friedman, a professor of clinical psychiatry and director of the psychopharmacology clinic at Weill Cornell Medical College in New York, they are simply molecular knockoffs of older ones. Some of the newer medications are somewhat safer, he says, but essentially, drug companies are just "tweaking the same molecules." Given the lack of ideas for effective new compounds and the high rate of failure for psychiatric drugs in expensive clinical trials—only about 8 percent succeed, compared with 15 percent for drugs overall—it's no wonder that, in Friedman's words, pharmaceutical companies have "gotten cold feet."

Indeed, in 2011 Novartis announced it was shutting down its center for basic neuroscience research in Basel, Switzerland. The company wasn't alone in stepping away from the search for psychiatric drugs. Over the last five years, other drug makers, including GlaxoSmithKline and AstraZeneca, have all scaled back efforts and decreased investments in neuroscience and related drug discovery. But Novartis's move was particularly noteworthy because the discovery of psychiatric drugs plays such a big part in its history. In the 1960s, Basel-based Sandoz, which merged with Ciba-Geigy in 1996 to form Novartis, was instrumental in developing clozapine, still one of the most effective treatments for severe schizophrenia. And Ciba, another Swiss parent of Novartis, had introduced imipramine in the late 1950s.

Now, in Cambridge, Novartis is back on the search. Ricardo Dolmetsch, the company's global head of neuroscience, is the one tasked with translating what he calls the recent revolutions in genetics and genomic tools into safe and effective medicines. A former professor of neuroscience at Stanford, Dolmetsch joined the company last August and immediately began hiring. Less than a year later, his colleagues are conducting experiments among stacks of plastic moving crates, even as they continue to set up the lab. Though there is a sense of excitement on his team, Dolmetsch's words are measured: "We now have the tools to give it another shot."

Jigsaw Puzzle

Faulty genes have a significant role in causing brain disorders. If you have an identical twin with schizophrenia, the likelihood that you will also have the disorder is between 40 and 65 percent; if a sibling has the illness, you have about a 10 percent chance. Statistics are similar for autism and bipolar disorder. While genes are somewhat less of a factor in depression than in the other disorders, they play a critical role there too. But, says Broad's Hyman, it's only in recent years that researchers have realized how complex the genetics are. When he was the director of NIMH in the 1990s, Hyman says, it was already clear to him and others that there was no single schizophrenia or autism gene. "But I thought that there were at most 20 or, at the extreme, 100 genes," he adds. "We were way off."

So far, researchers have identified hundreds of genetic variants associated with increased risk for schizophrenia, and Hyman guesses the number could go as high as a thousand. Some of the mutations appear to be common, while some rare variants seem to cause the same symptoms as those experienced by individuals with a completely different set of rare mutations. Moreover, different variants seem to confer different degrees of risk, and recent studies have shown that multiple disorders, including schizophrenia and autism, share a number of culpable genes. Hyman calls it an immensely complicated jigsaw puzzle.

Whether you think the extreme genetic complexity of brain disorders portends well for drug discovery depends on whether you're a pessimist or an optimist, says Pamela Sklar, who directs the department of psychiatric genomics at the Icahn School of Medicine at Mount Sinai in New York. The conventional approach to discovering drugs for diseases with a strong genetic component is to identify the gene causing or playing a prominent role in the illness, and then test compounds against the protein it codes for. That approach is not likely to work for most psychiatric illnesses, given that they are caused by combinations of so many genetic variants. But Sklar obviously leans toward optimism. She suggests that the numerous variants provide more chances to home in on key pathways involved in the disorders, and more opportunities to come up with clever ways to fix them.

The hope is that all those genetic variants will tend to affect common sets of molecular pathways, types of cells, or specific neurocircuits. That could help scientists pinpoint what's going wrong, and it could also give them new targets for potential treatments. Yet Sklar, who specializes in searching for the genetic causes of schizophrenia and bipolar disorder, acknowledges that despite the rapid advances in genetics over the last few years, large gaps in understanding remain. "We don't know all the risk factors," she says, "and with so few pieces of the puzzle, it's still hard to know how it all hangs together."

Add to this genetic mystery the fact that the brain has roughly 86 billion neurons and around a quadrillion synapses (the connecting points between neurons), and it becomes obvious how overwhelming it will be to understand the causes of brain disorders. It's why the ability to take cells from a patient and turn them into neurons in a dish has researchers so excited. Now they have a way to directly examine how genetic variants have affected the neurons of a patient with an illness. You still might not know all the details of the genetics, but at least you can see the results. What's more, new genome-editing methods make it possible to precisely alter the genes of the stem cells from which the neurons develop, adding a mutation associated with the illness to see how it affects the resulting neurons.

Researchers such as Stanford's Amit Etkin (below) and MIT's Kay Tye (right) believe improved understanding of neurocircuits and connections is critical to finding better treatments for psychiatric disorders.

LEAH FASTEN

But how do these neurons function in an actual brain, with its immense networks of circuits and connections? How are the genetic mutations implicated in autism and schizophrenia actually affecting those circuits to alter behavior? New research is starting to investigate those questions.

Later this summer, a colony of marmoset monkeys, primates native to South America, will begin living at MIT's McGovern Institute for Brain Research. Monkeys and people share a highly devel-

oped prefrontal cortex, the region near the front of the skull. And, says MIT's Feng, there's increasing evidence that many of the intractable defects of schizophrenia and the deficiencies of social communication and behavior found in autism can be traced to this area of the brain.

To begin to unravel what's going wrong, Feng and his colleagues plan to use genome editing to breed monkeys with precise mutations associated with psychiatric disorders. Initially, the scientists will focus on a rare mutation in a gene called *SHANK3*. Because it is an unusual example of a single gene that causes clear, autism-like changes in behavior, it's a simple place to start. Later generations of monkeys could have the multiple mutations found in most forms of autism and schizophrenia.

The monkeys could provide a more reliable way to test psychiatric drugs than rodents, whose brain circuitry is much less similar to ours. The idea is not to create animals with schizophrenia or autism—the complex mixture of aberrant human behaviors can't be truly replicated even in other primates—but rather to see how the genetic mutations change the circuitry at a molecular level and how the animals' behavior is altered as a result. "The behavior may not be identical to that in humans," says Feng, "but at least it's a readout. It becomes a confirmation that we can correct the circuits and that the changes lead to improvements in behavior."

How can this increasing knowledge of brain circuits and connections—and their role in such feelings as anxiety—be translated into actual therapies? Can researchers find effective and safe ways to intervene?

Don't Worry

The mouse cowers in one corner of the maze. Even in the video of the experiment, its anxiety seems palpable as it presses against a wall. A thin fiber-optic thread is connected to the animal's skull. Suddenly, after a burst of blue light through the fiber, the mouse begins scurrying about, exploring the four branches of the maze with new energy and courage.

The invention of optogenetics has revolutionized the study of neurocircuitry. But even among all the impressive studies using the technology, this mouse experiment, which Kay Tye conducted in 2011 as a postdoc at Stanford, stands out. Tye, now an assistant professor at MIT, showed she could turn anxiety on and off with a flip of a switch. Though she was targeting a part of the brain called the amygdala, which is well known to be involved in fear and anxiety, she was "surprised by just how sudden and robust the change was," she says: "It was almost instantaneous. I was blown away. It has changed how I think about the brain."

How can this increasing knowledge of brain circuits and connections—and their role in such feelings as anxiety—be translated into actual therapies? Can researchers find effective, safe ways to intervene in these circuits and connections in patients' brains—ways to fix what has gone wrong?

Optogenetics, at least in its existing version, doesn't look like the way to do that. It requires genetic modification of the cells that researchers want to activate and the intrusive use of a fiber-optic thread in the brain. That's why it's largely limited to rodents, though a few monkey experiments have been done. Using the technology to directly fix malfunctioning brain circuits in humans is not now practical, and it may never be. But as a research technique, it could give drug researchers what they so desperately need: new molecular targets. Researchers like Tye and Feng believe that their optogenetic experiments can help identify specific types of cells in the circuits underlying certain psychiatric symptoms. Then they will have to spot distinguishing markers on those cells that allow the drug to recognize them. It's an extremely promising approach; recent results do in fact suggest that there are ways to single out critical cells as targets for drug compounds. But the research is still in its very early days.

One alternative is to try to intervene directly in the circuits, skipping the use of drugs. A standard treatment for Parkinson's is to implant an electrode array in a patient's brain to calm tremors. It's called deep brain stimulation, and researchers at Emory University are attempting to adapt the technology to treat depression, inserting an electrode into a region of the brain called Area 25. Others are using deep brain stimulation to treat obsessive-compulsive disorder with promising results.

There might also be ways to directly affect faulty circuits without resorting to surgery. Amit Etkin, an assistant professor of psychiatry at Stanford, is using a combination of functional magnetic resonance imaging (fMRI) and noninvasive magnetic stimulation to map the circuitry that goes wrong in patients. His goal is to tailor magnetic stimulation, which is already widely used to treat intractable cases of depression, to the specific problems in a patient's neurocircuitry.

The therapy, which is administered using an electromagnetic coil placed against the scalp, uses magnetic pulses to create an electric current that can increase or decrease brain activity. The commercial version of the technique is designed to target the same small part of the prefrontal cortex in all patients, but by combining it with imaging technology, Etkin hopes to aim the stimulation more precisely to where a patient needs it. It won't be a miracle cure. The method seems to help only some people. But driving his work, says Etkin, is his frustration at not being able to offer patients more successful options.

Etkin, who also works at a clinic at the Palo Alto VA Hospital for veterans suffering from severe anxiety and depression, uses a variety of tools to help patients, including drugs and psychotherapy as well as magnetic stimulation. The key to making all the approaches more effective, he says, is to learn more about how faulty neural circuits and connections lead to aberrant behavior. In trying to fix those problems, he says, "I try not to be chauvinistic about the [treatment] technology." Working with patients not only motivated his research on potentially better treatments, he adds, but also teaches him about its practical limitations: "Lots of scientific studies might make sense, but there is a gap with what is doable in the real world. And sometimes the gap is significant."

Chasing Drugs

At Novartis, Ricardo Dolmetsch is responsible for trying to close that gap between the exploding scientific understanding of brain disorders and the availability of more effective drugs. And he's realistic about the prognosis: "I hope it will be a story of excitement. But we don't know yet. It takes a long time."

Dolmetsch is not your typical drug-industry manager. Less than a year ago, he was still running a lab at Stanford and helping to create a library of neurons from autism patients, to be hosted at the Allen Brain Institute in Seattle. And his Stanford website—he's officially taking a leave of absence from the university—still reflects his quirky persona. It has links to stories that describe his "impish humor" and his attempts to commute on a pogo stick during his early days at the university.

About a decade ago, his research took a dramatic turn. He had started out at Stanford looking into basic questions about the biochemistry in brain cells, work that was impressive enough to gain him an appointment as an assistant professor. But then, in 2005, his son received a diagnosis of autism. Frustrated by the lack of treatment options, Dolmetsch rebuilt his lab around researching the disorder. Since then, he has helped pioneer methods that take skin cells from individuals with autism, reprogram those cells to become stem cells, and then induce them to develop into neurons in which abnormalities can be studied. It's that technology, along with the revolution in genetics created in part by fast, cheap DNA sequencing, that he believes will be the linchpin of Novartis's revived effort to identify new psychiatric drugs.

"It allows you to start with the patient," says Dolmetsch. While oncologists have long had the ability to biopsy tumors, "you can't just drill holes in people's brains and take little bits out," he notes. "But now we have the capacity to make biopsies out of stem cells." Pushing the technology even further, the Novartis researchers are gearing up to make organoids, small chunks of brain that form in a petri dish as the neurons mature and lump together into three-dimensional structures. Not only can the researchers look for abnormalities, they can screen compounds in Novartis's vast library of potential drugs to see how the chemicals affect the neurons.

It's far from clear how or even whether the growing knowledge of what goes wrong in the brains of those with psychiatric disorders can lead to medicines. But at least now, researchers have the tools they need.

Designing drugs to precisely target circuits in the brain remains a more distant opportunity. "We now have some idea of the kinds of cells and regions of the brain that we need to inhibit or activate to make someone happier or less anxious. We no longer need to treat the brain as one mush of signaling molecules," Dolmetsch acknowledges. But that still leaves the daunting challenge of developing a drug that selectively activates or inactivates certain types of cells in certain circuits. "How do you do that?" he asks. "It's not been done before." He adds: "It's not where drug discovery is now, though it's where it is going."

Dolmetsch joined the pharmaceutical industry because he realized that the

Ricardo Dolmetsch hopes to resurrect psychiatric drug discovery at Novartis, using advances in genetics and stem-cell technologies.

science and technology had advanced far enough to create opportunities for developing new psychiatric drugs. He also realized after years of academic research that commercializing a new drug requires the resources, money, and patient populations available to a company like Novartis. It was, he says, "time to walk the walk" after years of talking about potential treatments for autism and schizophrenia.

Still, the failure to find effective new drugs for brain disorders—and the stigma that has grown around the high cost of those failures—is clearly never far from the minds of those in the industry. Companies "have developed some fantastic drugs over the last 10 years or so that were safe and engaged the target," says Dolmetsch, before delivering the punch line: the drugs were as ineffective as water. While he's convinced that he and his team have "a better way to discover new drugs," he also acknowledges that it will be five to eight years before they know whether their strategy based on the new genetic and stem-cell tools is working.

For those trying to discover new medicines, it's all about finding novel molecular targets that a chemical compound can safely affect in a way that will address the symptoms of a disease. That's a daunting challenge. And it's still far from clear how or even whether the growing knowledge of what goes wrong in the brains of those with psychiatric disorders can lead directly to such medicines. But at least now, after decades of dead ends, drug researchers finally have some of the tools they need to begin methodically testing strategies for finding and acting on those targets. When at last we have a better way to chase drugs for such disorders as autism and schizophrenia, says Dolmetsch, "it would be a bit of a crime not to take advantage of it." ■

David Rotman is the editor of MIT Technology Review.

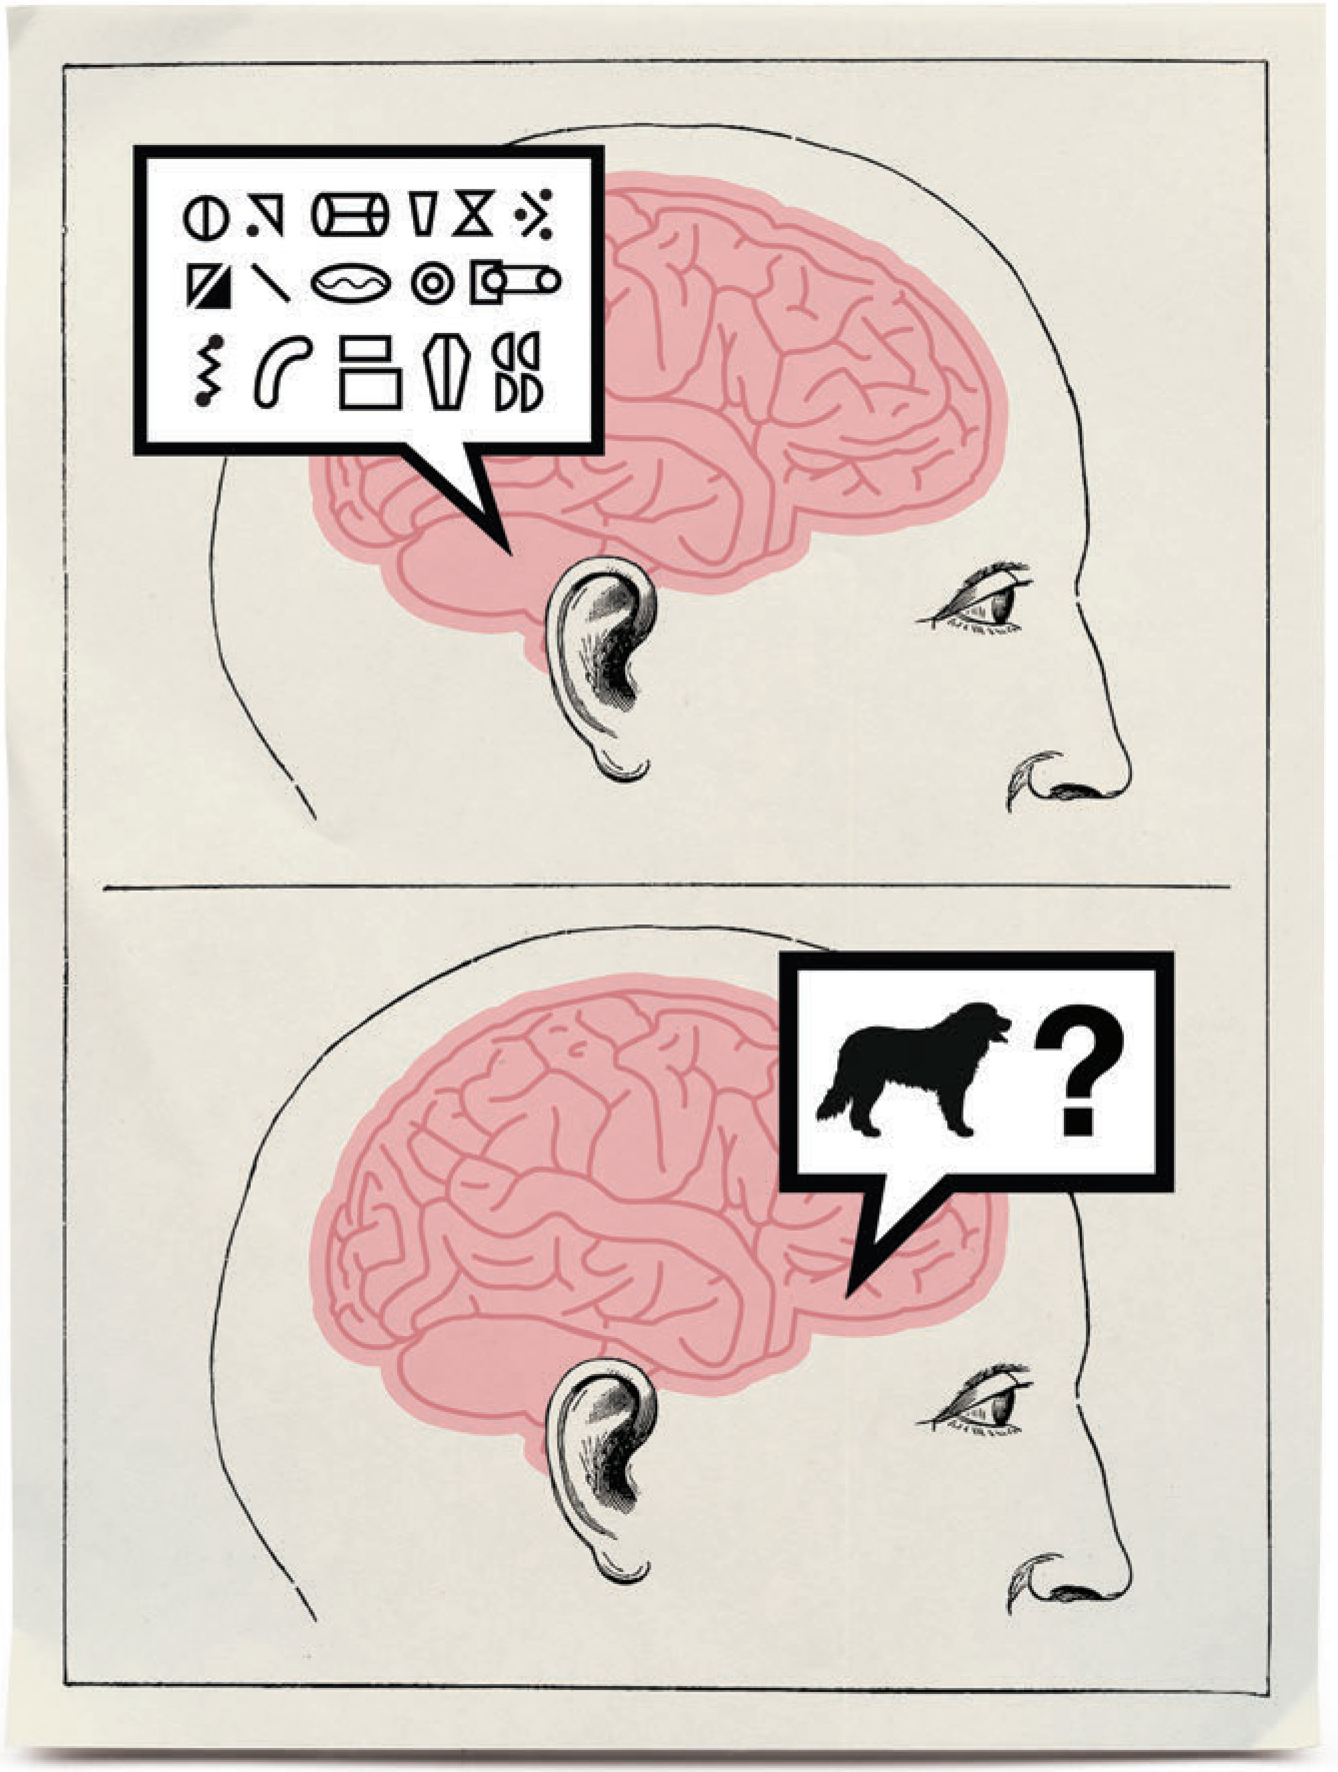

Matt Dorfman

Cracking the Brain's Codes

The brain relays information in complex ways based on the firing of neurons.

Decoding how those systems work could deepen our grasp of mental illness, lead to brain chips that enhance our memories, and enable computers to literally anticipate our needs.

By Christof Koch and Gary Marcus

In *What Is Life?* (1944), one of the fundamental questions the physicist Erwin Schrödinger posed was whether there was some sort of "hereditary code-script" embedded in chromosomes. A decade later, Crick and Watson answered Schrödinger's question in the affirmative. Genetic information was stored in the simple arrangement of nucleotides along long strings of DNA.

The question was what all those strings of DNA meant. As most schoolchildren now know, there was a code contained within: adjacent trios of nucleotides, so-called codons, are transcribed from DNA into transient sequences of RNA molecules, which are translated into the long chains of amino acids that we know as proteins. Cracking that code turned out to be a linchpin of virtually everything that followed in molecular biology. As it happens, the code for translating trios of nucleotides into amino acids (for example, the nucleotides AAG code for the amino acid lysine) turned out to be universal; cells in all organisms, large or small—bacteria, giant sequoias, dogs, and people—use the same code with minor variations. Will neuroscience ever discover something of similar beauty and power, a master code that allows us to interpret any pattern of neural activity at will?

At stake is virtually every radical advance in neuroscience that we might be able to imagine—brain implants that enhance our memories or treat mental disorders like schizophrenia and depression, for example, and neuroprosthetics that allow paralyzed patients to move their limbs. Because everything that you think, remember, and feel is encoded in your brain in some way, deciphering the activity of the brain will be a giant step toward the future of neuroengineering.

Someday, electronics implanted directly into the brain will enable patients with spinal-cord injury to bypass the affected nerves and control robots with their thoughts (see "The Thought Experiment," page 52). Future biofeedback systems may even be able to anticipate signs of mental disorder and head them off. Where people in the present use keyboards and touch screens, our descendants a hundred years hence may use direct brain-machine interfaces.

But to do that—to build software that can communicate directly with the brain—we need to crack its codes. We must learn how to look at sets of neurons, measure how they are firing, and reverse-engineer their message.

A Chaos of Codes

Already we're beginning to discover clues about how the brain's coding works. Perhaps the most fundamental: except in some of the tiniest creatures, such as the roundworm *C. elegans*, the basic unit of neuronal communication and coding is the spike (or action potential), an electrical impulse of about a tenth of a volt that lasts for a bit less than a millisecond. In the visual system, for example, rays of light entering the retina are promptly translated into spikes sent out on the optic nerve, the bundle of about one million output wires, called axons, that run from the eye to the rest of the brain. Literally everything that you see is based on these spikes, each retinal neuron firing at a different rate, depending on the nature of the stimulus, to yield several megabytes of visual information per second. The brain as a whole, throughout our waking lives, is a veritable symphony of neural spikes—perhaps one trillion per second. To a large degree, to decipher the brain is to infer the meaning of its spikes.

But the challenge is that spikes mean different things in different contexts. It is already clear that neuroscientists are unlikely to be as lucky as molecular biologists. Whereas the code converting nucleotides to amino acids is nearly universal, used in essentially the same way throughout the body and throughout the natural world, the spike-to-information code is likely to be a hodgepodge: not just one code but many, differing not only to some degree between different species but even between different parts of the brain. The brain has many functions, from controlling our muscles and voice to interpreting the sights, sounds, and smells that surround us, and each kind of problem necessitates its own kinds of codes.

A comparison with computer codes makes clear why this is to be expected. Consider the near-ubiquitous ASCII code representing the 128 characters, including numbers and alphanumeric text, used in communications such as plain-text e-mail. Almost every modern computer uses ASCII, which encodes the capital letter A as "100 0001," B as "100 0010," C as "100 0011," and so forth. When it comes to images, however, that code is useless, and different techniques must be used. Uncompressed bitmapped images, for example, assign strings of bytes to represent the intensities of the colors red, green, and blue for each pixel in the array making up an image. Different codes represent vector graphics, movies, or sound files.

Evidence points in the same direction for the brain. Rather than a single universal code spelling out what patterns of spikes mean, there appear to be many, depending on what kind of information is to be encoded. Sounds, for example, are inherently one-dimensional and vary rapidly across time, while the images that stream from the retina are two-dimensional and tend to change at a more deliberate pace. Olfaction, which depends on concentrations of hundreds of airborne odorants, relies on another system altogether. That said, there are some general principles. What matters most is not precisely when a particular neuron spikes but how often it does; the rate of firing is the main currency.

Consider, for example, neurons in the visual cortex, the area that receives impulses from the optic nerve via a relay in the thalamus. These neurons represent the world in terms of the basic elements making up any visual scene—lines, points, edges, and so on. A given neuron in the visual cortex might be stimulated most vigorously by vertical lines. As the line is rotated, the rate at which that neuron fires varies: four spikes in a tenth of a second if the line is vertical, but perhaps just once in the same interval if it is rotated 45° counterclockwise. Though the neuron responds most to vertical lines, it is never mute. No single spike signals whether it is responding to a vertical line or something else. Only in the aggregate—in the neuron's rate of firing over time—can the meaning of its activity be discerned.

This strategy, known as rate coding, is used in different ways in different brain systems, but it is common throughout the brain. Different subpopulations of neurons encode particular aspects of the world in a similar fashion—using firing rates to represent variations in brightness, speed, distance, orientation, color, pitch, and even haptic information like the position of a pinprick on the palm of your hand. Individual neurons fire most rapidly when they detect some preferred stimulus, less rapidly when they don't.

To make things more complicated, spikes emanating from different kinds of cells encode different kinds of information. The retina is an intricately layered piece of nervous-system tissue that lines the back of each eye. Its job is to transduce the shower of incoming photons into outgoing bursts of electrical spikes. Neuroanatomists have identified at least 60 different types of retinal neurons, each with its own specialized shape and function. The axons of 20 different retinal cell types make up the optic nerve, the eye's sole output. Some of these cells signal motion in several cardinal directions; others specialize in signaling overall image brightness or local contrast; still others carry information pertaining to color. Each of these populations streams its own data, in parallel, to different processing centers upstream from the eye. To reconstruct the nature of the information that the retina encodes, scientists must track not only the rate of every neuron's spiking but also the identity of each cell type. Four spikes coming from one type of cell may encode a small colored blob, whereas four spikes from a different cell type may encode a moving gray pattern. The number of spikes is meaningless unless we know what particular kind of cell they are coming from.

Some of the most important codes in any animal's brain are the ones it uses to pinpoint its location in space. How does our own internal GPS work? How do patterns of neural activity encode where we are as we move about?

And what is true of the retina seems to hold throughout the brain. All in all, there may be up to a thousand neuronal cell types in the human brain, each presumably with its own unique role.

Wisdom of Crowds

Typically, important codes in the brain involve the action of many neurons, not just one. The sight of a face, for instance, triggers activity in thousands of neurons in higher-order sectors of the visual cortex. Every cell responds somewhat differently, reacting to a different detail—the exact shape of the face, the hue of its skin, the direction in which the eyes are focused, and so on. The larger meaning inheres in the cells' collective response.

A major breakthrough in understanding this phenomenon, known as population coding, came in 1986, when Apostolos

Georgopoulos, Andrew Schwartz, and Ronald Kettner at the Johns Hopkins University School of Medicine learned how a set of neurons in the motor cortex of monkeys encoded the direction in which a monkey moves a limb. No one neuron fully determined where the limb would move, but information aggregated across a population of neurons did. By calculating a kind of weighted average of all the neurons that fired, Georgopoulos and his colleagues found, they could reliably and precisely infer the intended motion of the monkey's arm.

One of the first illustrations of what neurotechnology might someday achieve builds directly on this discovery. Brown University neuroscientist John Donoghue has leveraged the idea of population coding to build neural "decoders"—incorporating both software and electrodes—that interpret neural firing in real time. Donoghue's team implanted a brushlike array of microelectrodes directly into the motor cortex of a paralyzed patient to record neural activity as the patient imagined various types of motor activities. With the help of algorithms that interpreted these signals, the patient could use the results to control a robotic arm. The "mind" control of the robot arm is still slow and clumsy, akin to steering an out-of-alignment moving van. But the work is a powerful hint of what is to come as our capacity to decode the brain's activity improves.

Among the most important codes in any animal's brain are the ones it uses to pinpoint its location in space. How does our own internal GPS work? How do patterns of neural activity encode where we are? A first important hint came in the early 1970s with the discovery of what became known as place cells in the hippocampus of rats. Such cells fire every time the animal walks or runs through a particular part of a familiar environment. In the lab, one place cell might fire most often when the animal is near a maze's branch point; another might respond most actively when the animal is close to the entry point. The husband-and-wife team of Edward and May-Britt Moser discovered a second type of spatial coding based on what are known as grid cells. These neurons fire most actively when an animal is at the vertices of an imagined geometric grid representing its environment. With sets of such cells, the animal is able to triangulate its position, even in the dark. (There appear to be at least four different sets of these grid cells at different resolutions, allowing a fine degree of spatial representation.)

Other codes allow animals to control actions that take place over time. An example is the circuitry responsible for executing the motor sequences underlying singing in songbirds. Adult male finches sing to their female partners, each stereotyped song lasting but a few seconds. Neurons of one type within a particular structure are completely quiet until the bird begins to sing. Whenever the bird reaches a particular point in its song, these neurons suddenly erupt in a single burst of three to five tightly clustered spikes, only to fall silent again. Different neurons erupt at different times. It appears that individual clusters of neurons code for temporal order, each representing a specific moment in the bird's song.

Grandma Coding

Unlike a typewriter, in which a single key uniquely specifies each letter, the ASCII code uses multiple bits to determine a letter: it is an example of what computer scientists call a distributed code. In a similar way, theoreticians have often imagined that complex concepts might be bundles of individual "features"; the concept "Bernese mountain dog" might be represented by neurons that fire in response to notions such as "dog," "snow-loving," "friendly," "big," "brown and black," and so on, while many other neurons, such as those that respond to vehicles or cats, fail to fire. Collectively, this large population of neurons might represent a concept.

An alternative notion, called sparse coding, has received much less attention. Indeed, neuroscientists once scorned the idea as "grandmother-cell coding." The derisive term implied a hypothetical neuron that would fire only when its bearer saw or thought of his or her grandmother—surely, or so it seemed, a preposterous concept.

But recently, one of us (Koch) helped discover evidence for a variation on this theme. While there is no reason to think that a single neuron in your brain represents your grandmother, we now know that individual neurons (or at least comparatively small groups of them) can represent certain concepts with great specificity. Recordings from microelectrodes implanted deep inside the brains of epileptic patients revealed single neurons that responded to extremely specific stimuli, such as celebrities or familiar faces. One such cell, for instance, responded to different pictures of the actress Jennifer Aniston. Others responded to pictures of Luke Skywalker of *Star Wars* fame, or to his name spelled out. A familiar name may be represented by as few as a hundred and as many as a million neurons in the human hippocampus and neighboring regions.

Such findings suggest that the brain can indeed wire up small groups of neurons to encode important things it encounters over and over, a kind of neuronal shorthand that may be advantageous for quickly associating and integrating new facts with preëxisting knowledge.

Terra Incognita

If neuroscience has made real progress in figuring out how a given organism encodes what it experiences in a given moment, it has further to go toward understanding how organisms encode their long-term knowledge. We obviously wouldn't survive for long in this world if we couldn't learn new skills, like the orchestrated sequence of actions and decisions that go into driving a car. Yet the precise method by which we do this remains mysterious. Spikes are necessary but not sufficient for translating intention into action. Long-term memory—like the knowledge that we develop as we acquire a skill—is encoded differently, not by volleys of constantly circulating spikes but, rather, by literal rewiring of our neural networks.

That rewiring is accomplished at least in part by resculpting the synapses that connect neurons. We know that many different molecular processes are involved, but we still know little about which synapses are modified and when, and almost nothing about how to work backward from a neural connectivity diagram to the particular memories encoded.

Another mystery concerns how the brain represents phrases and sentences. Even if there is a small set of neurons defining a concept like your grandmother, it is unlikely that your brain has allocated specific sets of neurons to complex concepts that are less common but still immediately comprehensible, like "Barack Obama's maternal grandmother." It is similarly unlikely that the brain dedicates particular neurons full time to representing each new sentence we hear or produce. Instead, each time we interpret or produce a novel sentence, the brain probably integrates multiple neural populations, combining codes for basic elements (like individual words and concepts) into a system for representing complex, combinatorial wholes. As yet, we have no clue how this is accomplished.

One reason such questions about the brain's schemes for encoding information have proved so difficult to crack is that the human brain is so immensely complex, encompassing 86 billion neurons linked by something on the order of a quadrillion synaptic connections. Another is that our observational techniques remain crude. The most popular imaging tools for peering into the human brain do not have the spatial resolution to catch individual neurons in the act of firing. To study neural coding systems that are unique to humans, such as those used in language, we probably need tools that have not yet been invented, or at least substantially better ways of studying highly interspersed populations of individual neurons in the living brain.

It is also worth noting that what neuroengineers try to do is a bit like eavesdropping—tapping into the brain's own internal communications to try to figure out what they mean. Some of that eavesdropping may mislead us. Every neural code we can crack will tell us something about how the brain operates, but not every code we crack is something the brain itself makes direct use of. Some of them may be "epiphenomena"—accidental tics that, even if they prove useful for engineering and clinical applications, could be diversions on the road to a full understanding of the brain.

Nonetheless, there is reason to be optimistic that we are moving toward that understanding. Optogenetics now allows researchers to switch genetically identified classes of neurons on and off at will with colored beams of light. Any population of neurons that has a known, unique molecular zip code can be tagged with a fluorescent marker and then be either made to spike with millisecond precision or prevented from spiking. This allows neuroscientists to move from observing neuronal activity to delicately, transiently, and reversibly interfering with it. Optogenetics, now used primarily in flies and mice, will greatly speed up the search for neural codes. Instead of merely correlating spiking patterns with a behavior, experimentalists will be able to write in patterns of information and directly study the effects on the brain circuitry and behavior of live animals.

There is some cause for hope. Optogenetics now allows researchers to switch genetically identified classes of neurons on and off at will with colored beams of light. It could greatly speed up the search for codes.

Deciphering neural codes is only part of the battle. Cracking the brain's many codes won't tell us everything we want to know, any more than understanding ASCII codes can, by itself, tell us how a word processor works. Still, it is a vital prerequisite for building technologies that repair and enhance the brain.

Take, for example, new efforts to use optogenetics to remedy a form of blindness caused by degenerative disorders, such as retinitis pigmentosa, that attack the light-sensing cells of the eye. One promising strategy uses a virus injected into the eyeballs to genetically modify retinal ganglion cells so that they become responsive to light. A camera mounted on glasses would pulse beams of light into the retina and trigger electrical activity in the genetically modified cells, which would directly stimulate the next set of neurons in the signal path—restoring sight. But in order to make this work, scientists will have to learn the language of those neurons. As we learn to communicate with the brain in its own language, whole new worlds of possibilities may soon emerge. ■

Christof Koch is chief scientific officer of the Allen Institute for Brain Science in Seattle. Gary Marcus, a professor of psychology at New York University and a frequent blogger for the New Yorker, *is coeditor of the forthcoming book* The Future of the Brain.

Damasio has written highly acclaimed books that connect science to philosophy.

Brad Swonetz

For decades, biologists spurned emotion and feeling as uninteresting. But Antonio Damasio demonstrated that they were central to the life-regulating processes of almost all living creatures.

Damasio's essential insight is that feelings are "mental experiences of body states," which arise as the brain interprets emotions, themselves physical states arising from the body's responses to external stimuli. (The order of such events is: I am threatened, experience fear, and feel horror.) He has suggested that consciousness, whether the primitive "core consciousness" of animals or the "extended" self-conception of humans, requiring autobiographical memory, emerges from emotions and feelings.

His insight, dating back to the early 1990s, stemmed from the clinical study of brain lesions in patients unable to make good decisions because their emotions were impaired, but whose reason was otherwise unaffected—research made possible by the neuroanatomical studies of his wife and frequent coauthor, Hanna Damasio. Their work has always depended on advances in technology. More recently, tools such as functional neuroimaging, which measures the relationship between mental processes and activity in parts of the brain, have complemented the Damasios' use of neuroanatomy.

A professor of neuroscience at the University of Southern California, Damasio has written four artful books that explain his research to a broader audience and relate its discoveries to the abiding concerns of philosophy. He believes that neurobiological research has a distinctly philosophical purpose: "The scientist's voice need not be the mere record of life as it is," he wrote in a book on Descartes. "If only we want it, deeper knowledge of brain and mind will help achieve ... happiness."

Antonio Damasio talked with Jason Pontin, the editor in chief of *MIT Technology Review*.

Q+A
Antonio Damasio

JP: **When you were a young scientist in the late 1970s, emotion was not thought a proper field of inquiry.**

AD: We were told very often, "Well, you're going to be lost, because there's absolutely nothing there of consequence." We were pitied for our poor choice.

JP: **How so?**

AD: William James had tackled emotion richly and intelligently. But his ideas [mainly that emotions are the brain's mapping of body states, ideas that Damasio revived and experimentally verified] had led to huge controversies in the beginning of the 20th century that ended nowhere. Somehow researchers had the sense

"Mind begins at the level of feeling. It's when you have a feeling (even if you're a very little creature) that you begin to have a mind and a self. I'm ready to give the very teeny brain of an insect ... the possibility of having feelings. In fact, I would be flabbergasted to discover that they don't have feelings."

that emotion would not, in the end, be sufficiently distinctive—because animals had emotions, too. But what animals don't have, researchers told themselves, is language like we do, nor reason or creativity—so let's study *that*, they thought. And in fact, it's true that most creatures on the face of the earth do have something that could be called emotion, and something that could be called feeling. But that doesn't mean we humans don't use emotions and feelings in particular ways.

JP: **Because we have a conscious sense of self?**

AD: Yes. What's distinctive about humans is that we make use of fundamental processes of life regulation that include things like emotion and feeling, but we connect them with intellectual processes in such a way that we create a whole new world around us.

JP: **What made you so interested in emotions as an area of study?**

AD: There was something that appealed to me because of my interest in literature and music. It was a way of combining what was important to me with what I thought was going to be important scientifically.

JP: **What have you learned?**

AD: There are certain action programs that are obviously permanently installed in our organs and in our brains so that we can survive, flourish, procreate, and, eventually, die. This is the world of life regulation—homeostasis—that I am so interested in, and it covers a wide range of body states. There is an action program of thirst that leads you to seek water when you are dehydrated, but also an action program of fear when you are threatened. Once the action program is deployed and the brain has the possibility of mapping what has happened in the body, then *that* leads to the emergence of the mental state. During the action program of fear, a collection of things happen in my body that change me and make me behave in a certain way whether I want to or not. As that is happening to me, I have a mental representation of that body state as much as I have a mental representation of what frightened me.

JP: **And out of that "mapping" of something happening within the body comes a feeling, which is different from an emotion?**

AD: Exactly. For me, it's very important to separate emotion from feeling. We must separate the component that comes out of actions from the component that comes out of our perspective on those actions, which is feeling. Curiously, it's also where the self emerges, and consciousness itself. Mind begins at the level of feeling. It's when you have a feeling (even if you're a very little creature) that you begin to have a mind and a self.

JP: **But that would imply that only creatures with a fully formed sense of their minds could have fully formed feelings—**

AD: No, no, no. I'm ready to give the very teeny brain of an insect—provided it has the possibility of representing its body states—the possibility of having feelings. In fact, I would be flabbergasted to discover that they don't have feelings. Of course, what flies don't have is all the intellect around those feelings that could make use of them: to found a religious order, or develop an art form, or write a poem. They can't do that; but *we* can. In us, having feelings somehow allows us also to have creations that are responses to those feelings.

JP: **Do other animals have a kind of responsiveness to *their* feelings?**

AD: I'm not sure that I even understand your question.

JP: **Are dogs aware that they feel?**

AD: Of course. Of course dogs feel.

JP: **No, not "Do dogs feel?" I mean: is my dog Ferdinando conscious of feeling? Does he have feelings about his feelings?**

AD: *[Thinks.]* I don't know. I would have my doubts.

JP: **But humans are certainly conscious of being responsive.**

AD: Yes. We're aware of our feelings and are conscious of the pleasantness or unpleasantness associated with them. Look, what are the really powerful feelings that you deal with every day? Desires, appetites, hunger, thirst, pain—those are the basic things.

JP: **How much of the structure of civilization is devoted to controlling those basic things? Spinoza says that politics seeks to regulate such instincts for the common good.**

AD: We wouldn't have music, art, religion, science, technology, economics, politics, justice, or moral philosophy without the impelling force of feelings.

JP: **Do people emote in predictable ways regardless of their culture? For instance, does everyone hear the Western minor mode in music as sad?**

AD: We now know enough to say yes to that question.

At the Brain and Creativity Institute [which Damasio directs], we have been doing cross-cultural studies of emotion. At first we thought we would find very different patterns, especially with social emotions. In fact, we don't. Whether you are studying Chinese, Americans, or Iranians, you get very similar responses. There are lots of subtleties and lots of ways in which certain stimuli elicit different patterns of emotional response with different intensities, but the presence of sadness or joy is there with a uniformity that is strongly and beautifully human.

JP: **Could our emotions be augmented with implants or some other brain-interfacing technology?**

AD: Inasmuch as we can understand the neural processes behind any of these complex functions, once we do, the possibility of intervening is always there. Of course, we interface with brain function all the time: with diet, with alcohol, and with medications. So it's not that surgical interventions will be any great novelty. What will be novel is to make those interventions cleanly so that they are targeted. No, the more serious issue is the moral situations that might arise.

JP: **Why?**

AD: Because it really depends on what the intervention is aimed at achieving.

Suppose the intervention is aimed at resuscitating your lost ability to move a limb, or to see or hear. Do I have any moral problem? Of course not. But what if it interferes with states of the brain that are influential in how you make your decisions? Then you are entering a realm that should be reserved for the person alone.

JP: **What has been the most useful technology for understanding the biological basis of consciousness?**

AD: Imaging technologies have made a powerful contribution. At the same time, I'm painfully aware that they are limited in what they give us.

JP: **If you could wish into existence a better technology for observing the brain, what would it be?**

AD: I would not want to go to only one level, because I don't think the really interesting things occur at just one level. What we need are new techniques to understand the interrelation of levels. There are people who have spent a good part of their lives studying systems, which is the case with my wife and most of the people in our lab. We have done our work on neuroanatomy, and gone into cells only occasionally. But now we are actually studying the state of the functions of axons [nerve fibers in the brain], and we desperately need ways in which we can scale up from what we've found to higher and higher levels.

JP: **What would that technology look like?**

AD: I don't know. It needs to be invented. ■

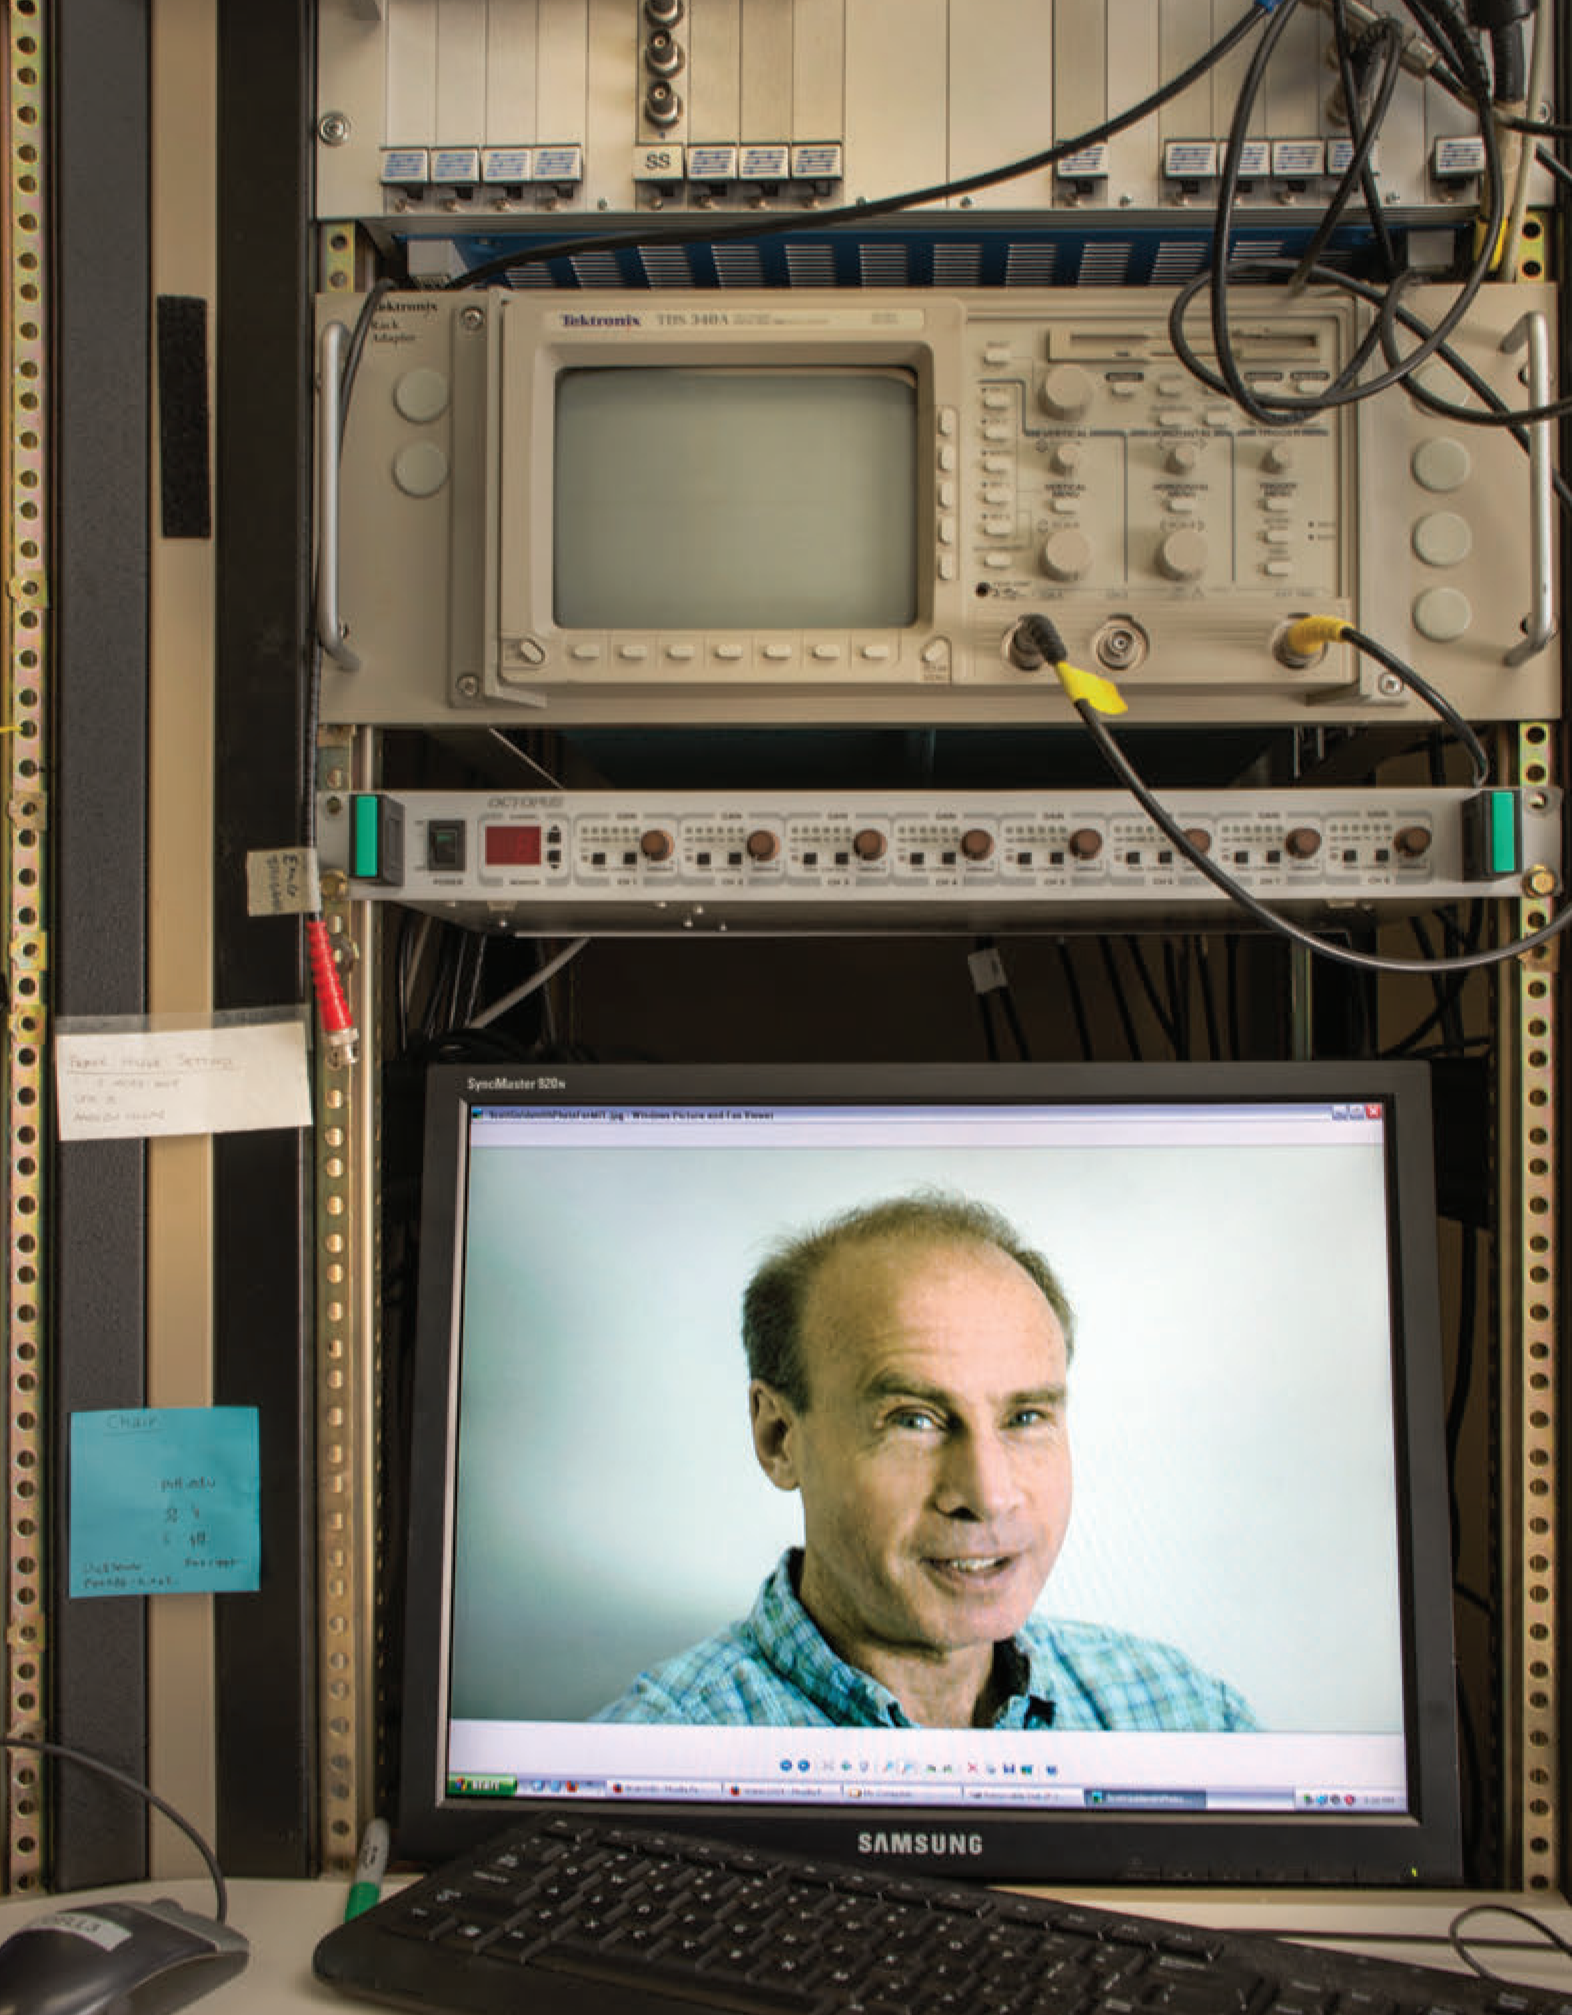
Tektronix
TDS 340A
SS
SyncMaster 920N
SAMSUNG

Neurophysiologist Andrew Schwartz at the laboratory where he connects brains to computers.

TH—E— —THOUGHT E—XP— E—R I—ME—NT

In a remarkable experiment, a paralyzed woman used her mind to control a robotic arm.

If only there were a realistic way to get this technology out of the lab and into real life. By Antonio Regalado

I WAS ABOUT 15 MINUTES LATE FOR MY first phone call with Jan Scheuermann. When I tried to apologize for keeping her waiting, she stopped me. "I wasn't just sitting around waiting for you, you know," she said, before catching herself. "Well, actually I *was* sitting around."

Scheuermann, who is 54, has been paralyzed for 14 years. She had been living in California and running a part-time business putting on mystery-theater dinners, where guests acted out roles she made up for them. "Perfectly healthy, married, with two kids," she says. One night, during a dinner she'd organized, it felt as if her legs were dragging behind her. "I chalked it up to being a cold snowy night, but there were a couple of steps in the house and boy, I was really having trouble," she says.

Anguished months of doctor's visits and misdiagnoses followed. A neurologist said she had multiple sclerosis. By then, she was using an electric wheelchair and "fading rapidly." She thought she was dying, so she moved home to Pittsburgh, where her family could take care of her children. Eventually she was diagnosed with a rare disease called spinocerebellar degeneration. She can feel her body, but the nerves that carry signals out from her brain no longer work. Her brain says "Move," but her limbs can't hear.

Two and a half years ago, doctors screwed two ports into Scheuermann's skull (she calls them "Lewis and Clark"). The ports allow researchers to insert cables that connect with two thumbtack-size implants in her brain's motor cortex.

Photographs by Scott Goldsmith

Jan Scheuermann stacks cones with a mind-controlled robot arm as research assistant Brian Wodlinger watches her work.

COURTESY OF UPMC

Two or three times a week, she joins a team of scientists at the University of Pittsburgh and is plugged into a robotic arm that she controls with her mind. She uses it to move blocks, stack cones, give high fives, and pose for silly pictures, doing things like pretending to knock out a researcher or two. She calls the arm Hector.

Scheuermann, who says that in her dreams she is not disabled, underwent brain surgery in 2012 after seeing a video of another paralyzed patient controlling a robotic arm with his thoughts. She immediately applied to join the study. During the surgery, doctors used an air gun to fire the two tiny beds of silicon needles, called the Utah Electrode Array, into her motor cortex, the slim strip of brain that runs over the top of the head to the jaws and controls voluntary movement. She awoke from the surgery with a pounding headache and "the worst case of buyer's remorse." She couldn't believe she'd had voluntary brain surgery. "I thought, *Please, God, don't let this be for nothing*. My greatest fear was that it wouldn't work," she says. But within days, she was controlling the robotic arm, and with unexpected success: "I was moving something in my environment for the first time in years. It was gasp-inducing and exciting. The researchers couldn't wipe the smile off their faces for weeks either."

Scheuermann is one of about 15 to 20 paralyzed patients who have joined long-term studies of implants that can convey information from the brain to a computer. She is the first subject at Pittsburgh. Nine others, including people in the advanced stages of ALS, have undergone similar tests in a closely related study, called BrainGate. Another four "locked-in" patients, unable to move or speak, have regained some ability to communicate thanks to a different type of electrode developed by a Georgia company called Neural Signals.

A third of these patients have undergone surgery since 2011, when the U.S. Food and Drug Administration said it would loosen rules for testing "truly pioneering technologies" such as brain-machine interfaces. More human experiments are under way. One, at Caltech, wants to give a patient "autonomous control over the Google Android tablet operating system." A team at Ohio State University, in collaboration with the R&D organization Battelle, put an implant in a patient in April with the intention of using the patient's brain signals to control stimulators attached to his arm. Battelle describes the idea as "reanimating a paralyzed limb under voluntary control by the participant's thoughts."

These nervy first-of-a-kind studies rely on the fact that recording the electrical activity of a few dozen cells in the brain can give a fairly accurate picture of where someone intends to move a limb. "We are technologically limited to sampling a couple of hundred neurons, from billions in your brain, so it's actually amazing they can get a signal out at all," says Kip Ludwig, director of the neural engineering program at the National Institute of Neurological Disorders and Stroke.

The technology being used at Pittsburgh was developed in physiology labs to study animals, and it is plainly still experimental. The bundled wires lead from Scheuermann's cranium to a bulky rack of signal processors, amplifiers, and computers. The nine-pound robotic arm, paid for by the military, has a dexterous hand and fingers that can make lifelike movements, but it is finicky, breaks frequently, and is somewhat dangerous. When things don't work, graduate students hunt among tangles of wires for loose connections.

John Donoghue, the Brown University neuroscientist who leads the longer-running BrainGate study, compares today's brain-machine interfaces to the first pacemakers. Those early models also featured carts of electronics, with wires punched through the skin into the heart. Some were hand-cranked. "When you don't know what is going on, you keep as much as possible on the outside and as little as possible on the inside," says Donoghue. Today, though, pacemakers are self-contained, powered by a long-lasting battery, and installed in a doctor's office. Donoghue says brain-machine interfaces are at the start of a similar trajectory.

like brushing teeth. During the study, Scheuermann's abilities were examined using the Action Research Arm Test, the same kit of wooden blocks, marbles, and cups that doctors use to evaluate hand dexterity in people with recent injuries. She scored 17 out of 57, or about as well as someone with a severe stroke. Without Hector, Scheuermann would have scored zero. The findings made *60 Minutes*.

Since the TV cameras went away, however, some of the shortcomings of the technology have become apparent. At first Scheuermann kept demonstrating new abilities. "It was success, success, success," she says. But controlling Hector has become harder. The reason is that the implants, over time, stop recording. The brain is a hostile environment for electronics, and tiny movements of the

At first it was "success, success, success," but Scheuermann says no one told her the implant might stop working. Gradually, it is recording from fewer neurons. Her control over the robot arm is weakening.

For brain-controlled computers to become a medical product, there has to be an economic rationale, and the risks must be offset by the reward. So far, Scheuermann's case has come closest to showing that these conditions can be met. In 2013, the Pittsburgh team reported its work with Scheuermann in the medical journal the *Lancet*. After two weeks, they reported, she could move the robot arm in three dimensions. Within a few months, she could make seven movements, including rotating Hector's hand and moving the thumb. At one point, she was filmed feeding herself a bite of a chocolate bar, a goal she had set for herself.

The researchers tried to show that they were close to something practical—helping with so-called "tasks of daily living" that most people take for granted,

array may build up scar tissue as well. The effect is well known to researchers and has been observed hundreds of times in animals. One by one, fewer neurons can be detected.

Scheuermann says no one told her. "The team said that they were expecting loss of neuron signals at some point. I was not, so I was surprised," she says. She now routinely controls the robot in only three to five dimensions, and she has gradually lost the ability to open and close its thumb and fingers. Was this at all like her experience of becoming paralyzed? I asked her the question a few days later by e-mail. She replied in a message typed by an aide who stays with her most days: "I was disappointed that I would probably never do better than I had already done, but accepted it without anger or bitterness."

Reanimation

The researcher who planned the Pittsburgh experiment is Andrew Schwartz, a lean Minnesotan whose laboratory occupies a sunlit floor dominated by three gray metal towers of equipment that are used to monitor monkeys in adjacent suites. Seen on closed-circuit TVs, the scenes from inside the experimental rooms defy belief. On one screen, a metal wheel repeatedly rotates, changing the position of a bright orange handle. After each revolution, an outsize robotic hand reaches up from the edge of the screen to grab the handle. Amid the spinning machinery, it's easy to miss the gray and pink face of the rhesus monkey that is controlling all this from a cable in its head.

The technology has its roots in the 1920s, with the discovery that neurons convey information via electrical "spikes" that can be recorded with a thin metal wire, or electrode. By 1969, a researcher named Eberhard Fetz had connected a single neuron in a monkey's brain to a dial the animal could see. The monkey, he discovered, learned to make the neuron fire faster to move the dial and get a reward of a banana-flavored pellet. Although Fetz didn't realize it at the time, he had created the first brain-machine interface.

Schwartz helped extend that discovery 30 years ago when physiologists began recording from many neurons in living animals. They found that although the entire motor cortex erupts in a blaze of electrical signals when an animal moves, a single neuron will tend to fire fastest in connection with certain movements—say, moving your arm left or up, or bending the elbow—but less quickly otherwise. Record from enough neurons and you can get a rough idea of what motion a person is making, or merely intending. "It's like a political poll, and the more neurons you poll, the better the result," he says.

The 192 electrodes on Scheuermann's two implants have recorded more than

The Utah Electrode Array, invented in the 1990s, has 96 silicon needles that record the electrical impulses of neurons inside the brain.

Five Milestones in Brain Control

1870 Scientists discover the motor cortex. They apply electricity to the brains of dogs, causing their limbs to move.

1969 At the University of Washington, monkeys are taught to move a dial using nerve impulses recorded from their brains. It is the first brain-machine interface.

1982 Electrical firing of neurons in the motor cortex is shown to predict which way a monkey's limb is moving. The discovery forms the basis for thought-controlled robots.

1998 Doctors implant a single electrode into the brain of a paralyzed man who cannot speak. He is able to move a cursor to select messages from a computer menu.

2014 Ohio doctors launch an attempt to "reanimate" a paralyzed man's limb by means of thought control. His brain signals will activate electrodes on his arm, making it move.

270 neurons at a time, which is the most ever simultaneously measured from a human being's brain. Schwartz says this is why her control over the robot has been so good.

The neuronal signals are interpreted by software called a decoder. Over the years, scientists built better and better decoders, and they tried more ambitious control schemes. In 1999, the Duke University neuroscientist Miguel Nicolelis trained a rat to swing a cantilever with its mind to obtain a reward. Three years later, Donoghue had a monkey moving a cursor in two dimensions across a computer screen, and by 2004 his BrainGate team had carried out the first long-term human test of the Utah array, showing that even someone whose limbs had been paralyzed for years could control a cursor mentally. By 2008, Schwartz had a monkey grabbing and feeding itself a marshmallow with a robotic hand.

Scheuermann has been able to quickly attempt many new tasks. She has been asked to control two robot arms at once and lift a box ("I only managed it once or twice," she says). Some results are strange: Scheuermann is able to close Hector's fingers around a plastic cone, but often only if she shuts her eyes first. Is the presence of the cone somehow reflected in the neurons' firing patterns? Schwartz has spent months trying to figure it out. Behind such points of uncertainty may lie major discoveries about how the brain prepares and executes actions.

Scheuermann once had her aide dress her in stick-on rat whiskers and a fuzzy tail to greet researchers. It was a darkly humorous way of acknowledging that these experiments depend on human volunteers. "They are not nearly as hard to train as these guys," Schwartz says, jerking a thumb to the row of monkey rooms.

These volunteers are trapped; some of them desperately hope science can provide an escape. Realistically, that is unlikely in their lifetimes. The first BrainGate volunteer was a 25-year-old named Matt Nagle, who had breathed through a ventilator since his spinal cord was severed during a knife fight. He was able to move a cursor on a screen in 2004. But Nagle also wanted to commit suicide and tried to get others to help him do it, according to *The Man with the Bionic Brain*, a book written by his doctor. He died of an infection in 2007. On online chat boards where paralyzed people trade hopeful news about possible cures, like stem cells, some dismiss brain-machine interfaces as wacky. Others are starting to think it's their best chance. "I'll take it! Cut off my dead arm and give me a robotic one that I can FEEL with please!" wrote one.

Some paralyzed people are starting to look at electronics as their best chance for recovery. "I'll take it! Cut off my dead arm and give me a robotic one that I can FEEL with please!" wrote one.

Schwartz says he hopes try to generate physical sensations from the robotic arm this year, if he can find another quadriplegic volunteer. Like Scheuermann, the next patient will receive two arrays in the motor cortex to control the robotic arm. But Schwartz says surgeons will place two additional implants into the volunteer's sensory cortex; these will receive signals from pressure sensors attached to the robotic fingertips. Studies by Nicolelis's Duke laboratory proved recently that animals do sense and respond to such electrical inputs. "We don't know if the subject

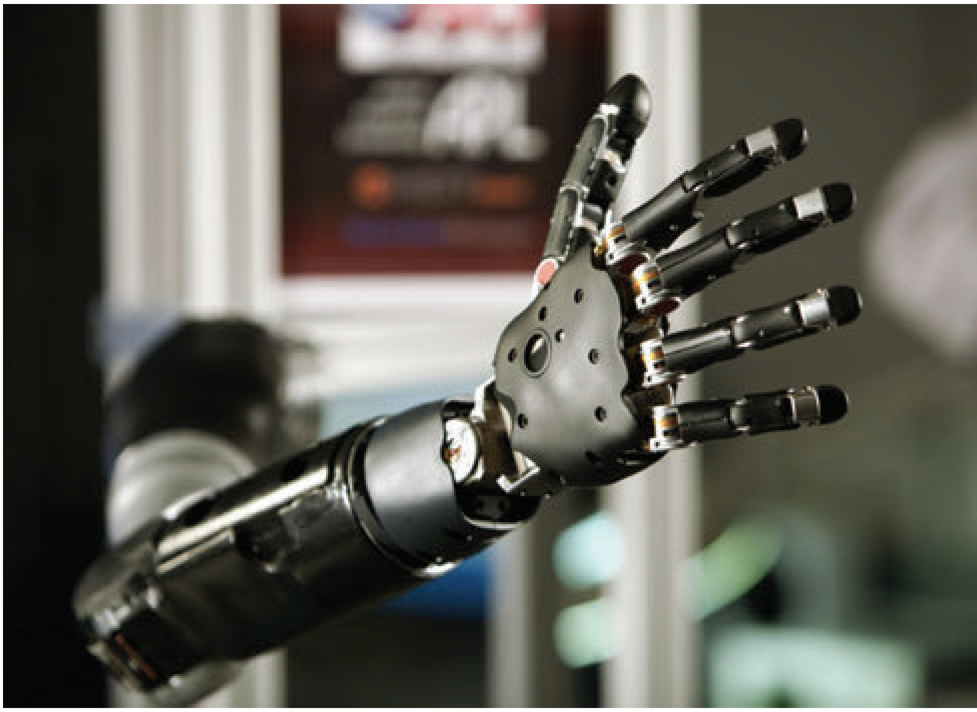

The Modular Prosthetic Limb was designed by Johns Hopkins University's Applied Physics Laboratory and funded by DARPA.

COURTESY OF DARPA AND JOHNS HOPKINS UNIVERSITY/APPLIED PHYSICS LAB

will feel it as touch," says Schwartz. "It's very crude and simplistic and an undoubtedly incorrect set of assumptions, but you can't ask the monkey what he just felt. We think it will be a new scientific finding. If the patient can say how it feels, that is going to be news."

Another key aim, shared by Schwartz and the BrainGate researchers, is to connect a volunteer's motor cortex to electrodes placed in his or her limbs, which would make the muscles contract—say, to open and close a hand. In April, the Ohio surgeons working with Battelle announced that they would be the first to try it. They put a brain implant in a man with a spinal-cord injury. And as soon as the patient recovers, says Battelle, they'll initiate tests to "reanimate" his fingers, wrist, and hand. "We want to help someone gain control over their own limb," says Chad Bouton, the engineer in charge of the project, who previously collaborated with the BrainGate group. "Can someone pick up a TV remote and change the channel?" Although Battelle has not won approval from regulators to attempt it, Bouton says the obvious next step is to try a bidirectional signal to and from a paralyzed limb, combining control and sensation.

Interface Problems

Brain-machine interfaces may seem as if they're progressing quickly. "If you fast-forward from the first video of that monkey to someone moving a robot in seven dimensions, picking things up, putting them down, it's pretty dramatic," says Lee Miller, a neurophysiologist at Northwestern University. "But what hasn't changed, literally, is the array. It's the Stanley Steamer of brain implants. Even if you demonstrate control, it's going to peter out in two to three years. We need an interface that will last 20 years before this can be a product."

The Utah array was developed in the early 1990s as a way to record from the cortex, initially of cats, with minimal trauma to the brain. It's believed that scar tissue builds up around the needle-like recording tips, each 1.5 millimeters long. If that interface problem is solved, says Miller, he doesn't see any reason why there couldn't be 100,000 people with brain implants to control wheelchairs, computer cursors, or their own limbs. Schwartz adds that if it's also possible to measure from enough neurons at once, someone could even play the piano with a thought-controlled robotic arm.

Researchers are pursuing several ideas for improving the brain interface. There are efforts to develop ultrathin electrodes, versions that are more compatible with the body, or sheets of flexible electronics that could wrap around the top of the brain. In San Francisco, doctors are studying whether such surface electrodes, although less precise, could be used in a decoder for speech, potentially allowing a person like Stephen Hawking to speak via a brain-computer interface. In an ambitious project launched last year at the University of California, Berkeley, researchers are trying to develop what they call "neural dust." The goal is to scatter micron-size piezoelectric sensors throughout the brain and use reflected sound waves to capture electrical discharges from nearby neurons.

Jose Carmena, a Berkeley researcher who, like Schwartz, works with monkeys to test the limits of thought control, now meets weekly with a group of a dozen scientists to outline plans for better ways to record from neurons. But whatever they come up with would have to be tested in animals for years before it could be tried in a person. "I don't think the Utah array is going to become the pacemaker of the brain," he says. "But maybe what we end up using is not that different. You don't see the newest computer in space missions. You need the most robust technology. It's the same kind of thing here."

Numbers Game

To succeed, any new medical device needs to be safe, useful, and economically viable. Right now, brain-machine interfaces don't meet these requirements. One problem is the riskiness of brain surgery and the chance of infection. At Brown, Donoghue says the BrainGate team is almost finished developing a wireless transmitter, about the size of a cigarette lighter, that would go under a person's skin and cut the infection risk by getting rid of the pedestals and wires that make brain-machine interfaces so unwieldy. Donoghue says that with a wireless system, implants could be a realistic medical option soon.

But that raises another tricky problem: what will patients control? The arm Scheuermann controls is still a hugely expensive prototype, and it often breaks. She worries that not everyone could afford one. Instead, Leigh Hochberg, a neurologist at Massachusetts General Hospital who runs the BrainGate study with Donoghue, thinks the first users will probably be "locked-in" patients who can neither move nor speak. Hochberg would consider it a "breakthrough" to afford such patients reliable thought control over a computer mouse. That would let them type out words or change the channel on a television.

Yet even locked-in patients can often move their eyes. This means they have simpler ways to communicate, like using an eye tracker. A survey of 61 ALS patients by the University of Michigan found that about 40 percent of them would consider undergoing surgery for a brain implant, but only if it would let them communicate more than 15 words a minute (a fifth of the people who responded to the survey were already unable to speak). BrainGate has not yet reported speeds that high.

All the pieces of the technology "have at some level been solved," says Andy Gotshalk, CEO of Blackrock Microsystems, which manufactures the Utah array and has acquired some of the BrainGate technology. "But if you ask me what is the product—is it a prosthetic arm or is it a wheelchair you control?—then I don't know. There is a high-level product in mind, which is to make life for quadriplegics a lot easier. But exactly what it would be hasn't been defined. It's just not concrete. The scientists are getting some high-level publications, but I have to think about the business plan, and that is a problem."

Without a clearly defined product to shoot for, no large company has jumped in to create a neuroprosthetic for quadriplegics. A startup went out of business after raising more than $30 million.

Without a clear product to shoot for, no large company has ever jumped in. And the risks for a business are especially high because there are relatively few patients with complete quadriplegia—about 40,000 in the U.S.—and even fewer with advanced ALS. A company Donoghue created, Cyberkinetics, went out of business after raising more than $30 million. Researchers instead get by on small grants that are insignificant compared with a typical commercial effort to develop a new medical device, which can cost $100 million. "There is not a single company out there willing to put the money in to create a neuroprosthetic for quadriplegics, and the market is not big enough for a venture capitalist to get in," says Gotshalk. "The numbers don't add up."

Others think the technology behind brain-machine interfaces may have unexpected applications, far removed from controlling robot arms. Many researchers, including Carmena and the team at Battelle, are trying to determine whether the interfaces might help rehabilitate stroke patients. Because they form a large market, it "would be a game changer," Carmena says. Some of the recording technologies could be useful for understanding psychiatric diseases like depression or obsessive-compulsive disorder.

In Scheuermann's case, at least, her brain-machine interface has proved to be powerful medicine. When she first arrived at Pittsburgh, her doctors say, her affect was flat, and she didn't smile. But being part of the experiment energized her. "I was loving it. I had coworkers for the first time in 20 years, and I felt needed," she says. She finished dictating a mystery novel, *Sharp as a Cucumber*, that she'd started before she became ill and published it online. Now she's working on a second one. Scheuermann told me she'd like to have a robotic arm at home. She'd be able to open the door, roll onto her lawn, and talk to neighbors. Maybe she'd open the refrigerator and grab a sandwich that her aide had prepared for her.

Our call was ending. The moment was awkward. I could hang up the phone, but she couldn't. Her husband had gone out shopping. Hector was back in the lab. She was alone and couldn't move. "That's all right," Scheuermann said. "I'll just let the phone slip to the floor. Good-bye."

Antonio Regalado is MIT Technology Review*'s senior editor for business.*

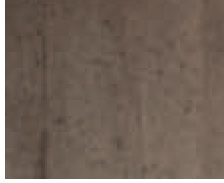

The ability to discern what other people are thinking and feeling is critical to social interaction and a key part of the human experience. So it's not surprising that the human brain devotes a lot of resources to so-called social cognition. But only recently has neuroscience begun to tease apart which brain regions and processes are devoted to thinking about other people.

Understanding how the brain perceives, interprets, and makes decisions about other people could help advance treatments and interventions for autism and other disorders in which social interactions are impaired. It could also help us build more socially intelligent computers. So far, artificial intelligence has struggled to program computers that make the kinds of social judgments that come easily to us, such as interpreting ambiguous facial expressions or deciding whether a person's words convey anger or sadness.

More than a decade ago, neuroscientist Rebecca Saxe discovered a brain region that develops a "theory of mind"—a sense of what other people are thinking and feeling. More recently, she became an investigator in MIT's Center for Minds, Brains, and Machines, and she's been studying autism and social cognition in children and adults. Saxe and *MIT Technology Review* contributing editor Courtney Humphries discussed the implications of this new research on the social brain.

Q+A

Rebecca Saxe

CH: **Is social cognition something seen only in humans?**

RS: There's every reason to believe that at least in some ways, we are uniquely good at this kind of thing. Humans are by far the most social species, other than insects. Even the interaction you and I are having, that two strangers could meet and for no particular reason act coöperatively for an hour—that's unheard of outside of humans. If two ants did it, they would be sisters. Our extraordinary social lives and our hugely complex cognitive capacities combine to make human social cognition distinctive.

CH: **How do you study it in the brain?**

RS: Nothing invasive—no genetic engineering, no optogenetics, none of these things. We are limited to what are called noninvasive neuroimaging technologies, and the most well known of these is functional MRI, which uses blood flow as an index of neural activity.

CH: **So you can see which areas of the brain are active when people are thinking about other people. Was it a surprise to find brain regions devoted to social cognition?**

RS: In a sense, that had been predicted about 15 or 20 years earlier, when people noticed that kids with autism seemed disproportionately bad at that kind of thing. But otherwise this was completely unknown. In some ways, I think it was the most important, most surprising new discovery of human cognitive neuroscience. All the visual

Saxe studies how theory of mind develops and which systems in the brain support it.

Jared Leeds

regions, all the sensory regions, all the motor control regions—we predicted they would be there. But the social brain was not predicted at all. It just emerged. That was wild.

In the last 10 years [we've been] trying to refine our interpretation about what information is in those brain regions, how they interact with one another, how they develop, whether or not those brain regions do have anything to do with autism.

CH: And do these regions indeed not function well in people with autism?

RS: That was the original hypothesis that we went after. Maybe [people with autism] are trying to solve social problems with the machinery we would use for other problems, instead of having the dedicated machinery. There is no evidence that that is right. Too bad, because I like that idea. Autism has turned out to be a much, much harder problem at every level of analysis than I think anyone expected. Ten years ago people thought that cognitively, neurally, genetically, autism would be crackable. Now it looks like maybe there are thousands of genetic variations of autism.

CH: How might your work help lead to more socially capable computers?

RS: To me, the signature of human social cognition is the same thing that makes good old-fashioned AI hard, which is its generativity. We can recognize and think about and reason through a literally infinite set of situations and goals and human minds. And yet we have a very particular and finite machinery to do that. So what are the right ingredients? If we know what those are, then we can try to understand how the combinations of those ingredients generate this massively productive, infinitely generalizable human capacity.

CH: What do you mean by "ingredients"?

RS: Let's say you hear about a friend of yours. She was told she was being called to her boss's office, and she thought she was finally getting the promotion she'd been waiting for. But it turned out she actually got fired. Let's say the next day you see her coming down the street and she has a huge smile on her face. Probably not what you had expected, right?

You take that and you build a whole interior world. Maybe it's a fake smile and she's putting on a brave face. Maybe she's relieved because now she can move to the other side of the continent and live with her boyfriend. You need to figure out: What were her goals? What did she want? What changed her mind? There are all kinds of features of that story that you were able to extract in the moment.

If a computer could extract [such] features, we could [improve its ability to do] sentiment analysis. There's a huge focus in AI right now on trying to take the natural language people use and figure out: Did they like or not like that thing? Did they like that restaurant or not like that restaurant? Now take it up to the level of distinguishing between language when you feel disappointed, lonely, or terrified. That's the kind of problem that we want to solve.

> **"There's a huge focus in AI right now on trying to take the natural language people use and figure out [whether they like something]. Now take it up to the level of distinguishing between language when you feel disappointed, lonely, or terrified. That's the kind of problem that we want to solve."**

CH: How can computers learn to do that?

RS: You need to translate those words into more abstract things—goals, desires, plans. My colleague Josh Tenenbaum and I have been working for years just to build a kind of mathematical representation of what it means to think of somebody as having a plan or a goal, such that this model can predict human judgments about the person's goal in a really simple context. What do you need to know about a goal? We're trying to build models that describe that knowledge.

CH: That's very different from having a computer look at millions of examples to find patterns.

RS: Exactly. This is not big data; it's trying to describe the structure of the knowledge. That's always been viewed as an opposition: the people who want bigger data sets and the people who want the right knowledge structures. My impression right now is that there's a lot more intermediate ground. What used to be viewed as opposite traditions in AI should now be viewed as complementary, where you try to figure out probabilistic representations that learn from data.

CH: But the prospect of replicating social cognition in a computer seems far off, right? We don't yet understand how the brain does it.

RS: It feels pretty plausible that the full understanding is not in the grasp of me in my lifetime, and that's good, because it means I have a lot of work to do. So in the meantime, I do whatever seems likely to produce a little bit of instrumental progress toward that bigger goal. ■

Leonard Greco

MIT TECHNOLOGY REVIEW
VOL. 117 | NO. 4

Searching for the "Free Will" Neuron

Gabriel Kreiman's single-neuron measurements of unconscious decision-making may not topple Descartes, but they could someday point to ways we can learn to control ourselves.

By
David Talbot

IT WAS AN EXPEDITION SEEKING SOMEthing never caught before: a single human neuron lighting up to create an urge, albeit for the minor task of moving an index finger, before the subject was even aware of feeling anything. Four years ago, Itzhak Fried, a neurosurgeon at the University of California, Los Angeles, slipped several probes, each with eight hairlike electrodes able to record from single neurons, into the brains of epilepsy patients. (The patients were undergoing surgery to diagnose the source of severe seizures and had agreed to participate in experiments during the process.) Probes in place, the patients—who were conscious—were given instructions to press a button at any time of their choosing, but also to report when they'd first felt the urge to do so.

Later, Gabriel Kreiman, a neuroscientist at Harvard Medical School and Children's Hospital in Boston, captured the quarry. Poring over data after surgeries in 12 patients, he found telltale flashes of individual neurons in the pre-supplementary motor area (associated with movement) and the anterior cingulate (associated with motivation and attention), preceding the reported urges by anywhere from hundreds of milliseconds to several seconds. It was a direct neural measurement of the unconscious brain at work—caught in the act of formulating a volitional, or freely willed, decision. Now Kreiman and his colleagues are planning to repeat the feat, but this time they aim to detect pre-urge signatures in real time and stop the subject from performing the action—or see if that's even possible.

A variety of imaging studies in humans have revealed that brain activity related to decision-making tends to precede conscious action. Implants in macaques and other animals have examined brain circuits involved in perception and action. But Kreiman broke ground by directly measuring a preconscious decision in humans at the level of single neurons. To be sure, the readouts came from an average of just 20 neurons in each patient. (The human brain has about 86 billion of them, each with thousands of connections.) And ultimately, those neurons fired only in response to a chain of even earlier events. But as more such experiments peer deeper into the labyrinth of neural activity behind decisions—whether they involve moving a finger or opting to buy, eat, or kill something—science could eventually tease out the full circuitry of decision-making and perhaps point to behavioral therapies or treatments. "We need to understand the neuronal basis of voluntary decision-making—or 'freely willed' decision-making—and its pathological counterparts if we want to help people such as drug, sex, food, and gambling addicts, or patients with obsessive-compulsive disorder," says Christof Koch, chief scientist at the Allen Institute of Brain Science in Seattle (see "Cracking the Brain's Codes," page 42). "Many of these people perfectly well know that what they are doing is dysfunctional but feel powerless to prevent themselves from engaging in these behaviors."

Kreiman, 42, believes his work challenges important Western philosophical ideas about free will. The Argentine-born neuroscientist, an associate professor at Harvard Medical School, specializes in visual object recognition and memory formation, which draw partly on unconscious processes. He has a thick mop of black hair and a tendency to pause and think a long moment before reframing a question and replying to it expansively. At the wheel of his Jeep as we drove down Broadway in Cambridge, Massachusetts, Kreiman leaned over to adjust the MP3 player—toggling between Vivaldi, Lady Gaga, and Bach. As he did so, his left hand, the one on the steering

wheel, slipped to let the Jeep drift a bit over the double yellow lines. Kreiman's view is that his neurons made him do it, and they also made him correct his small error an instant later; in short, all actions are the result of neural computations and nothing more. "I am interested in a basic age-old question," he says. "Are decisions really free? I have a somewhat extreme view of this—that there is nothing really free about free will. Ultimately, there are neurons that obey the laws of physics and mathematics. It's fine if you say 'I decided'—that's the language we use. But there is no god in the machine—only neurons that are firing."

Our philosophical ideas about free will date back to Aristotle and were systematized by René Descartes, who argued that humans possess a God-given "mind," separate from our material bodies, that endows us with the capacity to freely choose one thing rather than another. Kreiman takes this as his departure point. But he's not arguing that we lack any control over ourselves. He doesn't say that our decisions aren't influenced by evolution, experiences, societal norms, sensations, and perceived consequences. "All of these external influences are fundamental to the way we decide what we do," he says. "We do have experiences, we do learn, we can change our behavior."

But the firing of a neuron that guides us one way or another is ultimately like the toss of a coin, Kreiman insists. "The rules that govern our decisions are similar to the rules that govern whether a coin will land one way or the other. Ultimately there is physics; it is chaotic in both cases, but at the end of the day, nobody will argue the coin 'wanted' to land heads or tails. There is no real volition to the coin."

Testing Free Will

It's only in the past three to four decades that imaging tools and probes have been able to measure what actually happens in the brain. A key research milestone was reached in the early 1980s when Benjamin Libet, a researcher in the physiology department at the University of California, San Francisco, made a remarkable study that tested the idea of conscious free will with actual data.

Libet fitted subjects with EEGs—gadgets that measure aggregate electrical brain activity through the scalp—and had them look at a clock dial that spun around every 2.8 seconds. The subjects were asked to press a button whenever they chose to do so—but told they should also take note of where the time hand was when they first felt the "wish or urge." It turns out that the actual brain activity involved in the action began 300 milliseconds, on average, before the subject was conscious of wanting to press the button. While some scientists criticized the methods—questioning, among other things, the accuracy of the subjects' self-reporting—the study set others thinking about how to investigate the same questions. Since then, functional magnetic resonance imaging (fMRI) has been used to map brain activity by measuring blood flow, and other studies have also measured brain activity processes that take place before decisions are made. But while fMRI transformed brain science, it was still only an indirect tool, providing very low spatial resolution and averaging data from millions of neurons. Kreiman's own study design was the same as Libet's, with the important addition of the direct single-neuron measurement.

When Libet was in his prime, Kreiman was a boy. As a student of physical chemistry at the University of Buenos Aires, he was interested in neurons and brains. When he went for his PhD at Caltech, his passion solidified under his advisor, Koch. Koch was deep in collaboration with Francis Crick, co-discoverer of DNA's structure, to look for evidence of how consciousness was represented by neurons. For the star-struck kid from Argentina, "it was really life-changing," he recalls. "Several decades ago, people said this was not a question serious scientists should be thinking about; they either had to be smoking something or have a Nobel Prize"—and Crick, of course, was a Nobelist. Crick hypothesized that studying how the brain processed visual information was one way to study consciousness (we tap unconscious processes to quickly decipher scenes and objects), and he collaborated with Koch on a number of important studies. Kreiman was inspired by the work. "I was very excited about the possibility of asking what seems to be the most fundamental aspect of cognition, consciousness, and free will in a reductionist way—in terms of neurons and circuits of neurons," he says.

One thing was in short supply: humans willing to have scientists cut open their skulls and poke at their brains. One day in the late 1990s, Kreiman attended a journal club—a kind of book club for scientists reviewing the latest literature—and came across a paper by Fried on how to do brain science in people getting electrodes implanted in their brains to identify the source of severe epileptic seizures. Before he'd heard of Fried, "I thought examining the activity of neurons was the domain of monkeys and rats and cats, not humans," Kreiman says. Crick introduced Koch to Fried, and soon Koch, Fried, and Kreiman were collaborating on studies that investigated human neural activity, including the experiment that made the direct neural measurement of the urge to move a finger. "This was the opening shot in a new phase of the investigation of questions of voluntary action and free will," Koch says.

Better Decisions

A perennial debate in philosophy is whether, if our choices are caused by something (anything), we can still be said to pos-

BRUCE PETERSON

Inserted into the brain, electrodes such as this one detect neural activity.

sess free will. Hilary Bok, a philosopher at Johns Hopkins University, says many modern philosophers—perhaps most—believe freedom of decision is possible, but that of course neural processes lead to urges and actions. "The idea that your choice might be caused—including by something happening in the brain—occurred to us a long time before the neuroscientists started filling in the details of how it happened," she says. Freedom doesn't require a ghost in the machine; we might still have some sort of free will if it can be shown that our neural circuits give us the capacity to weigh options and choose the right ones. "I love these experiments and think they are really interesting," she says, "but I'm less convinced whether they have shown anything crucial about free will."

What's really important about the experiments, she adds, is that they start to provide insights into human behavior. This could someday lead to therapies, but until then, insight alone can help. Consider the case of James Fallon, a neuroscientist at the University of California, Irvine, who discovered that his own fMRI scan bore similarities to those of known psychopaths (it indicated low activity in brain regions associated with self-control and empathy). Fallon has described how he now makes a conscious effort to modify his everyday decisions and behaviors—such as his tendency to want to defeat his young grandchildren at games. "When I think about freedom of the will, a part of what is required is that we have some ability to control our own actions," Bok says. "It would be important to me to discover whether a psychopath who decides to use his or her narcissism to defeat that narcissism can actually succeed."

Though it's still at an early stage, Kreiman's work adds to this kind of understanding, says Patricia Churchland, a philosopher at the University of California, San Diego, who avers that neuroscience can illuminate old philosophical questions. Churchland believes that the experiments could be important in shedding light on whether decisions can be altered or urges held in check, and might help explain why some people have difficulty controlling their impulses after brain damage. "The explanations are far from complete, but a pattern of results is fitting together in illuminating ways," she says. "Self-control is an entirely real brain phenomenon. Insofar as self-control is a key component of free choice, we do in fact have free choice. From a range of data it is becoming quite clear that there are significant neural differences between people who have the capacity to cancel actions or defer gratification and those whose capacity is diminished."

There is no "god in the machine"—only neurons that are firing, Kreiman says.

Kreiman, too, sees practical potential in teasing apart the circuitry of decision-making but deflects my questions about how his work could lead to new drugs or therapies. "The main question at the scientific level is to understand the mechanism by which volitional decisions are made: where, when, and how they are orchestrated," he says.

In his quest for a better picture of how that orchestration works, Kreiman now has a new collaborator in Ed Boyden, a neuroscientist at MIT who has developed novel tools for analyzing brain circuits. Among other efforts, Boyden is testing far denser neural probes in mice; these have the potential to record from as many as 100 times more neurons simultaneously (see "Neuroscience's New Toolbox," page 20). This technology could allow scientists to identify many of the neurons involved in making an "urge" neuron fire. "If we can get that, it would be transformative for this project as well as many others," Kreiman says. Projects around the world that are currently working to map the brain's circuitry would especially stand to benefit.

With such tools, rather than just seeing a single neuron light up, you might see a network of electrical signals that made it happen. Then you could, perhaps, see what lit the neuron that makes you wiggle the finger, and the neuron that makes you reach for the bottle. "If you can map the neural activity and see how neurons are dynamically generating the outcomes—it's like seeing how the brain is computing a decision," Boyden says. "You'd like to be able to see how the emotions, sensations, and memories work together."

While Kreiman sees no free will, he does believe mechanisms of self-control are built into the circuits that guide him down Broadway and through life. He wants to discover them, but he concedes that even if he were to do so today, "tonight, everything will be the same." It may be that the illusion of free will is part of the wiring—and impossible to shake off. ■

David Talbot is chief correspondent of MIT Technology Review.

MIT Technology Review
BUSINESS REPORT

The Internet of Things

Billions of computers that can sense and communicate from anywhere are coming online. What will it mean for business?

CONTENTS

Find additional material for this report at technologyreview.com/business

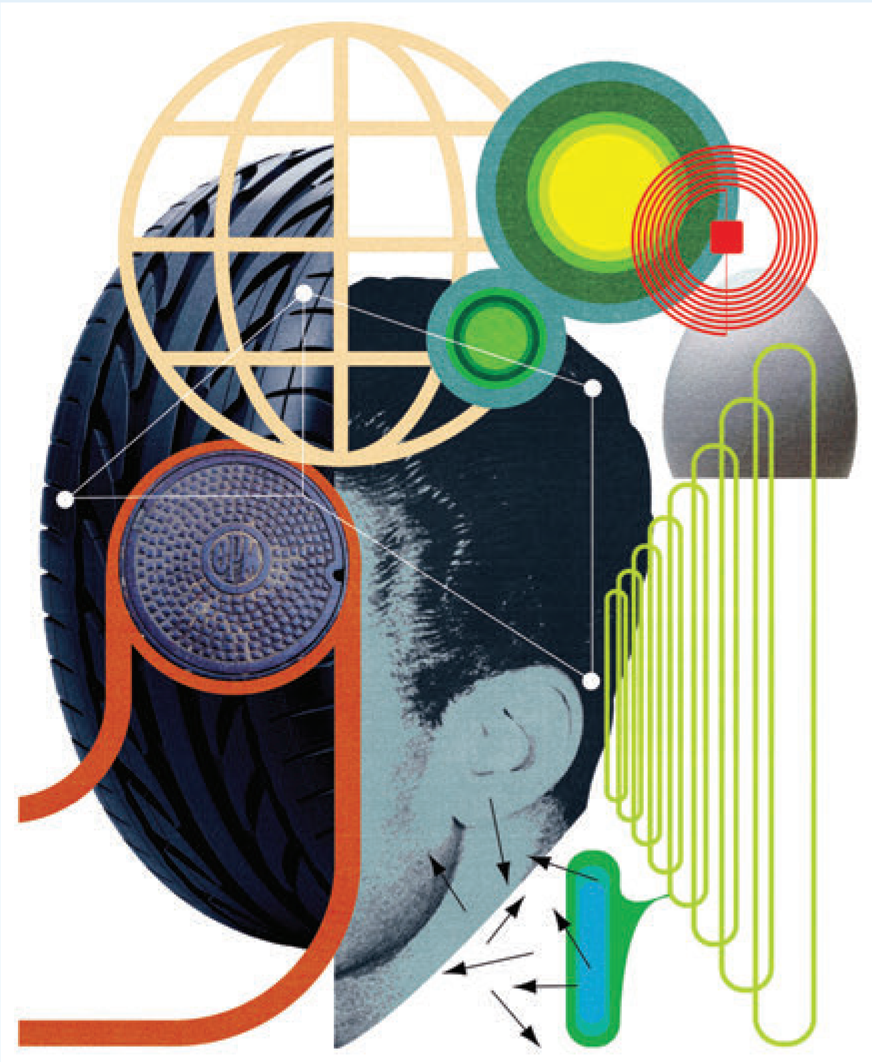

STUART BRADFORD

The Big Question

Business Adapts to a New Style of Computer

Are companies ready for billions of everyday objects to join the Internet?

● The technology industry is preparing for the Internet of things, a type of computing characterized by small, often dumb, usually unseen computers attached to objects. These devices sense and transmit data about the environment or offer new means of controlling it.

For more than a decade technologists have predicted and argued about the onslaught of these ubiquitous devices. "There is a lot of quibbling about what to call it, but there's little doubt that we're seeing the inklings of a new class of computer," says David Blaauw, who leads a lab at the University of Michigan that makes functioning computers no bigger than a typed letter o.

A key feature is very cheap radios, etched right into silicon. There's one in your smartphone. But now prices are falling to around $5. As they get →

There’s a way
to do it better.
Find it. –Thomas Edison

NATIONAL INSTRUMENTS™

Each year, over 35,000 companies choose a better approach to measurement and control.

Machines Go Online

The number of everyday objects, or "things," connecting to the Internet will exceed PCs and smartphones.

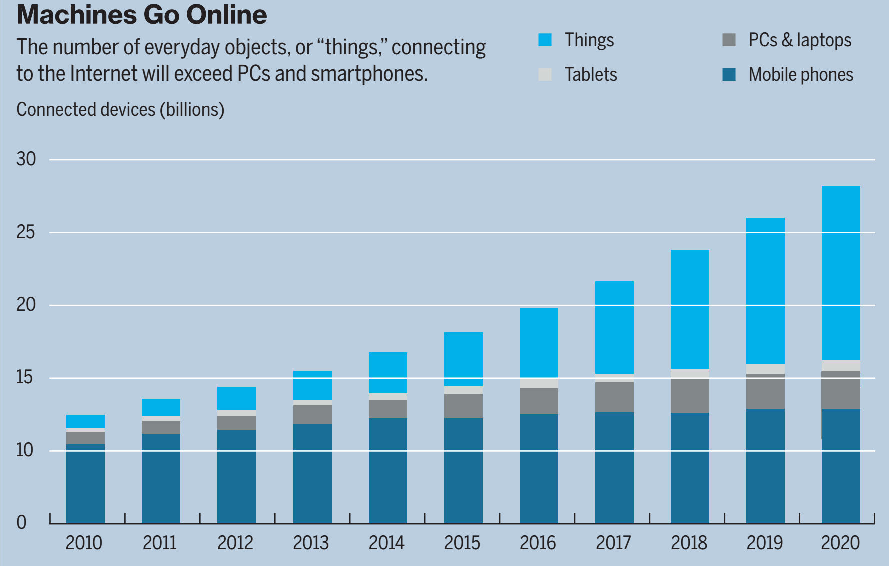

SOURCE: GSMA

As computers with wireless capability become cheap, it's becoming affordable to connect more things to the Internet, like sensors in sewer pipes, factory machinery, lights, and home appliances.

cheaper, it's becoming affordable to connect more things, like sewer pipes or trash cans. At the University of California, Berkeley, researchers are even designing computers the size of a pinhead to collect data inside the brain and transmit it through the skull. The idea is that human bodies will join the network, too.

It can all sound far-fetched and overhyped. Does anyone really need a smart coffee pot or a refrigerator with a Web browser? Plenty of the inventions do seem silly. On Amazon, product reviewers have had a field day with a $78 digital "egg minder" that reports to a smartphone which egg in a refrigerator is oldest. "Wonderful product!" sneered one. "So many gray hairs avoided by never having to worry about my eggs again."

Yet for every killer app that wasn't, there's another computer-sensor combination that has quietly added to the capabilities of some machine. Since 2007, for instance, every new car in the United States has had a chip in each tire that measures pressure and sends data by radio to the car's central computer. It's starting to add up. The average new car has 60 microprocessors in it, according to the Center for Automotive Research. Electronics account for 40 percent of the cost of making a car.

The Internet of things is especially important for companies that sell network equipment, like Cisco Systems. Cisco has been enthusiastically predicting that 50 billion "things" could be connected to communications networks within six years, up from around 10 billion mobile phones and PCs today. Another beneficiary is the $300 billion semiconductor industry. As Blaauw notes, "Every time there has been a new class of computing, the total revenue for that class was larger than the previous ones. If that trend holds, it means the Internet of things will be bigger yet again."

60
Number of microprocessors in the average new car

But every shift promises pain, too. Large companies like Intel are already reeling from the rapid emergence of smartphones. Intel, with its powerful, power-hungry chips, was shut out of phones. So was Microsoft. Now both these companies, and many others, are groping to find the winning combination of software, interfaces, and processors for whatever comes next.

And it's not just technology companies that must stay alert this time around. The reason, explains Marshall Van Alstyne, a professor at Boston University, is that as ordinary products become connected, their manufacturers may enter information businesses whose economics are alien to them. It's one thing to manufacture shoes, but what about a shoe that communicates? Products could turn out to be valuable mainly as the basis for new services. "You might find the data is more valuable than the shoe," says Van Alstyne.

In this *MIT Technology Review* business report we decided to explore the big question of what new businesses will arise as things get connected. One company making the point is Nest Labs, maker of a slick-looking smart thermostat that's coupled to the Internet. Nest, which was acquired by Google this year, has been clobbering rival thermostat makers. But now that it has a network of thermostats and can control them from afar, it's starting to offer services to electric utilities. On hot days it can selectively turn down air conditioners, controlling demand.

Nest's tests with utilities are still small. But one day, with a few bits sent across a network, the company might put a power plant or two out of business. No wonder this year, in his annual letter to shareholders, Jeff Immelt, CEO of General Electric, the world's largest manufacturer, told his investors that "every industrial company will be a software company."

Gordon Bell, a Microsoft researcher and a pioneer of the original computer revolution, believes no one knows exactly what form computing will take on the Internet of things. But he says that's unsurprising. The importance of the PC and the smartphone became clear only after their development. "The 'Internet of things' is a way of saying that more of the world will become part of the network," he says. "That is what is going on. We are assimilating the world into the computer. It's just more and more computers." *—Antonio Regalado*

Leaders

The Economics of the Internet of Things

As everyday objects get connected, brace yourself for network effects, says one economist.

● Product companies compete by building ever bigger factories to turn out ever cheaper widgets. But a very different sort of economics comes into play when those widgets start to communicate. It's called the network effect—when each new user of a product makes its value higher. Think of the telephone a century ago. The greater the number of people who used Bell's invention, the more valuable it became to all of them. The telephone became a platform for countless new businesses its inventor never imagined.

Now that more objects are getting wired up into networks—street lights, wind turbines, automobiles—there are opportunities for new platforms to emerge. That's why some companies are seeking the advice of Marshall Van Alstyne, a business professor at Boston University who has studied the economics of e-mail spam and social networks.

These days, Van Alstyne studies "platform economics," or why companies such as Uber, Apple, and Amazon are so successful—and what traditional product makers can do to emulate them. *MIT Technology Review*'s senior editor for business, Antonio Regalado, visited Van Alstyne at his office in Boston.

How can I tell if a business is a platform?
If you produce the value, then you are a classic product company. But there are new systems where value is being created outside the firm, and that's a platform business. Apple gets 30 percent of the cut from other people's innovations in its app store. I define a platform as a published standard that lets others connect to it, together with a governance model, which is the rules of who gets what. Business platforms are often engaged in consummating a match. It's a match between riders and drivers with Uber. It's between travelers and spare capacity of guest rooms in Airbnb.

Is connecting ordinary objects, like toasters, to the Internet going to trigger new platforms?
Absolutely, yes. But you can't stop at the connectivity. The technologist's mistake is often to stop simply at the standards, the connections. You also have to add the reasons for other people to add value. That often means allowing recombination of features in ways that you, the original designer, just cannot anticipate. People have combined the functions of the iPhone into hundreds of thousands of apps that Apple never even conceived of. That is also what the Internet of things enables if you design it in the right way.

What's an example of this happening?
Philips Lighting just called me. They are adding a series of APIs to their LED lights so anyone can create millions of colors, create romantic mood apps or the colors of a sunset from one of your favorite trips. You can change the lights in your study in conjunction with the stock market conditions. That is the Internet of things, and they're opening it to anyone.

Do product companies have a difficult time making this kind of transition?
They have a really difficult time with the mental models. It's fascinating. Most companies compete by adding new features to products. They haven't been in the business of thinking of how to add new communities or network effects. One of the points I make is that platform business models are like playing 3-D chess.

You estimate that half the top 20 companies in the world, like Google, own platforms. Why are they winning?
There is a strong argument that platforms beat products every time. Think of how the iPhone is absorbing the features of the voice recorder, the calculator, and game consoles. The reason for this is that as a stand-alone product, you're going to have a certain pace of innovation. But if you have opened your product so that third parties can add value, and you have designed the rules of the ecosystem such that they want to, your innovation curve is going to be faster.

To me this means there are huge opportunities to take away business from existing players in all different kinds of goods. Or for existing players to expand their markets if they are paying attention.


COURTESY OF REDACTIE EMERCE

"Most companies compete by adding new features to products. They haven't been in the business of thinking of how to add new communities or network effects."
—Marshall Van Alstyne

What are some of the next areas for platforms?
It's where you see connectivity is coming in. Cities, health care, education, electricity grids.

What are the biggest challenges?
In many cases, the governance models have not been established. For instance, population density can be determined by mobile-phone distribution. A telecom company owns that data. How do you motivate them to share it? All these sensors are capturing data, but how do you divide the value? Those are the rules that need to be worked out, and that's the missing piece of most of these discussions about the Internet of things. You have to build economic incentives around it, not simply connectivity.

Peak Power

On a 104° day in Austin, remote control of home thermostats helped cut power demand.

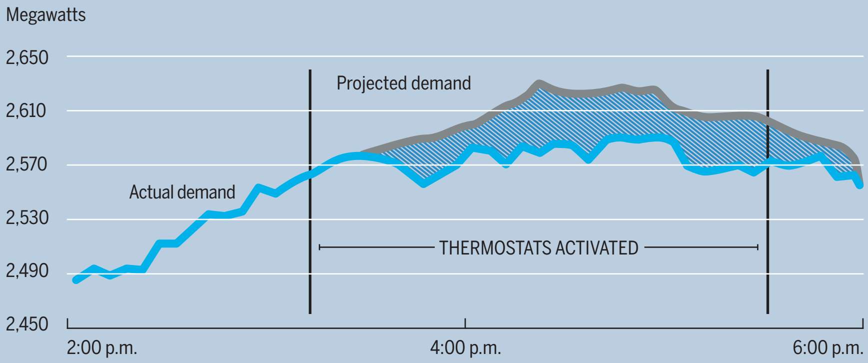

Wired Home

TVs, heaters, and other appliances will account for more of the Internet-connected devices in the average U.S. home.

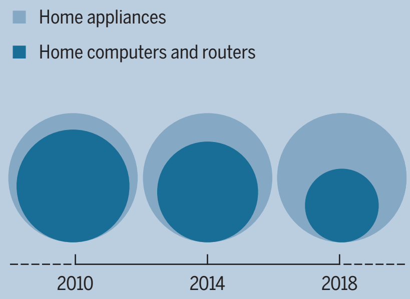

Case Studies

The Lowly Thermostat, Now Minter of Megawatts

How Nest is turning its consumer hit into a service for utilities.

● Google's $3.2 billion acquisition of Nest Labs in January put the Internet of things on the map. Everyone had vaguely understood that connecting everyday objects to the Internet could be a big deal. Here was an eye-popping price tag to prove it.

Nest, founded by former Apple engineers in 2010, had managed to turn the humble thermostat into a slick, Internet-connected gadget. By this year, Nest was selling 100,000 of them a month, according to an estimate by Morgan Stanley.

At $249 a pop, that's a nice business. But more interesting is what Nest has been up to since last May in Texas, where an Austin utility is paying Nest to remotely turn down people's air conditioners in order to conserve power on hot summer days—just when electricity is most expensive.

For utilities, this kind of "demand response" has long been seen as a killer app for a smart electrical grid, because if electricity use can be lowered just enough at peak times, utilities can avoid firing up costly (and dirty) backup plants.

Demand response is a neat trick. The Nest thermostat manages it by combining two things that are typically separate—price information and control over demand. It's consumers who control the air conditioners, electric heaters, and furnaces that dominate a home's energy diet. But the actual cost of energy can vary widely, in ways that consumers only dimly appreciate and can't influence.

While utilities frequently carry out demand response with commercial customers, consumers until now have shown little interest. Nest Labs' breakthrough was to make a device that has popular appeal. "There's a lot of digital Internet thermostats out there, but Nest was able to create a concept around it. They've created something that people are relating to," says Mary Ann Piette, a demand response expert and head of the Building Technology and Urban Systems Department at Lawrence Berkeley National Laboratory.

Once inside a home, Nest starts its real work: gathering data. It has motion detectors; sensors for temperature, humidity, and light; and algorithms that learn residents' habits and preferences and can program heating and AC settings. A Wi-Fi connection brings in weather data and allows consumers to control the system with a phone or Web browser.

Data is just the start. Just as Google parlays what it knows about you into tools for advertisers on the Web, Nest is using its capabilities to create new types of services for utilities to buy. "We can go to utilities and say, 'We've actually got a lot of customers in your service territory who already have a Nest,'" says Scott McGaraghan, Nest Labs' head of energy products. "And [then we] can flip it on."

Austin's municipal utility, Austin Energy, is one of five utilities that have signed up for Nest Labs' Rush Hour Rewards, as the service is called. Air conditioners account for half of Texas's electricity demand on hot days, and that demand for cooling drives the wholesale cost of electricity from less than $40 per megawatt-hour to well over $1,000.

Twelve months ago Austin Energy started offering a one-time $85 rebate to customers who agreed to let it automatically trim their air-conditioning using smart thermostats sold by Nest and other companies. Each company earns $25 for

Once inside a home, Nest starts its real work: gathering data. It has a motion detector; sensors for temperature, humidity, and light; and algorithms that learn residents' habits and preferences.

SOURCE: ANALYSIS MASON; SOURCE: AUSTIN ENERGY; COURTESY OF NEST every thermostat it enrolls, and another $15 per thermostat each year after that.

The "vast majority" of the 5,500 thermostats registered so far are Nests, according to Sarah Talkington, the Austin Energy engineer leading the program. Nest says it finds that roughly half its customers will sign up for demand response when the opportunity is offered.

By the end of last summer, Talkington says, she could log on to a Nest portal and, with a few keystrokes, dial down the next day's demand by nearly 5.7 megawatts. That may seem small compared with the 2,800 megawatts that often sizzle across the Austin grid, but every watt counts. On hot days like September 3, 2013, as temperatures rose to 104 °F, the cost of power spiked to a record $4,900 per megawatt-hour.

Austin had tried residential demand response before, using one-way pagers to turn air conditioners on and off. But the utility couldn't know if customers were home, so it wasn't able to shut off any one air conditioner for long. Nest, in contrast, builds a thermal model of each house and predicts how quickly it will warm up. It can also guess whether people will be home. The result, says McGaraghan, is that Nest can maximize energy savings and minimize annoyance to residents.

Talkington predicts the residential program will enroll enough homes to save more than 13 megawatts through demand response this summer. Even if Austin gives out $2 million in rebates, that is cheaper than increasing power supply by building a natural-gas-fired generator. According to Michael Webber, co-director of the clean-energy incubator at the University of Texas in Austin, new power supply costs $500,000 to $4,000,000 per megawatt of capacity, depending on the type of plant.

Webber believes that within five years the "vast preponderance" of Texans will have smart thermostats. And Nest knows that whoever builds this network first could win big, especially as other energy-consuming devices, like electric cars and hot-water heaters, also get wired up.

Eventually, the effects of demand response could be profound. Austin's program is designed to manage demand only during the 50 hours each year when electricity consumption tests the grid's limits most. But if demand response can expand to cover the 300 or 400 hours of peak usage, it could entirely shut down the market for "peakers," or gas-fired plants that come online only to sell expensive electricity. "That's a big chunk of money that's at stake," says Tom Osterhus, CEO of Integral Analytics, a Cincinnati-based maker of smart-grid analytics software. "It's in the billions." *—Peter Fairley*

Case Studies

The Light Bulb Gets a Digital Makeover

Electric lights are 135 years old. The Internet is 45. They're finally getting connected.

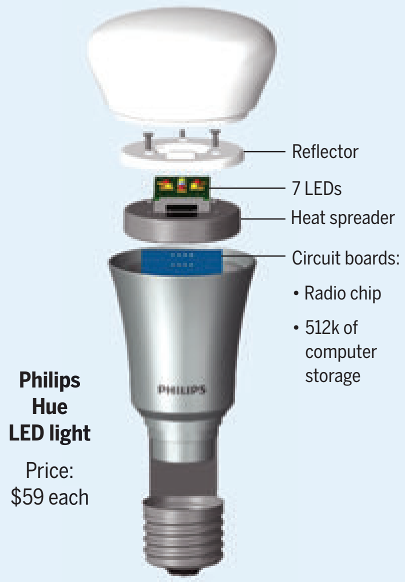

COURTESY OF PHILIPS

● To demonstrate how the Internet is changing one of the oldest and least exciting technology businesses around, Shane De Lima, an engineer at Philips Lighting, took out his smartphone. A flick across the screen sent a message to a nearby Wi-Fi router and then to a wireless hub, which shot a radio command to a chip in the base of an LED lamp in front of us.

A moment later, the conference room where we were sitting darkened.

It may seem like Rube Goldberg's idea of how to turn off a light. Or it could be the beginning of how lighting companies such as Philips find their way from selling lighting hardware into networks, software, apps, and new kinds of services.

The introduction of networked lights is happening because of another trend. Manufacturers have been replacing incandescent and fluorescent lights with ultra-efficient LEDs, or light-emitting diodes. The U.S. Department of Energy says that LEDs had 4 percent of the U.S. lighting market in 2013, but it predicts this figure will rise to 74 percent of all lights by 2030.

Because LEDs are solid-state devices that emit light from a semiconductor chip, they already sit on a circuit board. That means they can readily share space with sensors, wireless chips, and a small computer, allowing light fixtures to become networked sensor hubs.

For example, last year Philips gave outside developers access to the software that runs its Hue line of residential LED lights. Now it's possible to download Goldee, a smartphone app that turns your house the color of a Paris sunset, or Ambify, a $2.99 app created by a German programmer that makes the lights flash to music as in a jukebox.

That's a very different kind of business from selling light bulbs, as Philips has done since 1891. "With the new digitization of light, we have only begun to scratch the surface on how we can control it, integrate it with other systems, and collect rich data," says Brian Bernstein, Philips's global head of indoor lighting systems.

Another look at how lighting systems are changing will emerge this November, when a 14-story regional headquarters for Deloitte, nearing completion in Amsterdam, will be festooned with networked LEDs in each fixture—the first such installation for Philips.

Each of 6,500 light fixtures will have an IP address and five sensors—all of them wired only to Ethernet cables. (They'll use "power over Ethernet" technology to deliver the juice to each fixture as well as data.) The fixtures include a light →

sensor to dim the LEDs during the day, and a motion detector that covers the area directly beneath each light and turns the light off when no one is there. "We expect to spend 70 percent less on light, because systems [give] us much more control," says Erik Ubels, chief information officer at Deloitte in the Netherlands. Additional sensors in the LED fixtures can monitor temperature, humidity, carbon dioxide, and heat, turning the lights into a kind of building-management system.

Prices for LEDs are high but falling quickly. A "dumb" LED that puts out as much light as a $1.25 incandescent bulb now sells for $9 (but uses one-sixth the energy and lasts much longer). That's down from $40 each a couple of years ago. A connected LED bulb from Philips's Hue line retails in the U.S. for $59. But these will get cheaper, too. Philips says a third of its lighting revenue now comes from LEDs, and about 1.7 percent from the newer LEDs that can connect to the Internet.

Many other uses are being explored. A department store in Dusseldorf, Germany, is using LEDs to send out light frequencies that communicate with shoppers' smartphones. Philips has placed street lights in Barcelona that react to how many people are strolling by. *—David Talbot*

Leaders

GE's $1 Billion Software Bet

To protect lucrative business servicing machines, GE turns to industrial Internet.

● To understand why General Electric is plowing $1 billion into the idea of using software to transform industry, put yourself in the shoes of Jeff Immelt, its CEO.

As recently as 2004, GE had reigned as the most valuable company on the planet. But these days, it's not even the largest in America. Apple, Microsoft, and Google are all bigger. Software is king of the hill. And, as Immelt came to realize, GE is not that great at software.

Service Play

GE's industrial revenues are split between selling products and servicing them.

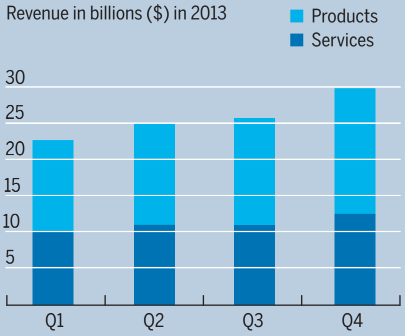

Internal surveys had discovered that GE sold $4 billion worth of industrial software a year—the kind used to run pumps or monitor wind turbines. That's as much as the total revenue of Salesforce.com. But these efforts were scattered and not always state-of-the-art. And that gap was turning dangerous. GE had always believed that since it knew the materials and the physics of its jet engines and medical scanners, no one could best it in understanding those machines. But companies that specialize in analytics, like IBM, were increasingly spooking GE by figuring out when big-ticket machines like a gas turbine might fail—just by studying raw feeds from gauges or vibration monitors.

This was no small thing. GE sells $60 billion a year in industrial equipment. But its most lucrative business is servicing the machines. Now software companies were looking to take a part of that pie, to get between GE and its largest source of profits. As Immelt would later say, "We cannot afford to concede how the data gathered in our industry is used by other companies."

In 2012, GE unveiled its answer to these threats, a campaign it calls the "industrial Internet." It included a new research lab across the bay from Silicon Valley, where it has hired 800 people, many of them programmers and data scientists.

"People have told companies like GE for years that they can't be in the software business," Immelt said last year. "We're too slow. We're big and dopey. But you know what? We are extremely dedicated to winning in the markets we're in. And this is a to-the-death fight to remain relevant to our customers."

Peter Evans, then a GE executive, was given the job of shaping what he calls the "meta-narrative" around GE's big launch. Industrial companies, which prize reliability, aren't nearly as quick to jump for new

"This is a to-the-death fight to remain relevant to our customers."

—General Electric CEO Jeffrey Immelt

technology as consumers. So GE's industrial-Internet pitch was structured around the huge economic gains even a 1 percent improvement in efficiency might bring to a number of industries if they used more analytics software. That number was fairly arbitrary—something safe, "just 1 percent," recalls Evans. But here Immelt's marketing skills came into play. "Not 'just 1 percent,'" he said, flipping it around. GE's slogan would be "The Power of 1 Percent."

In a stroke, GE had shifted the discussion about where the Internet was going next. Other companies had been talking about connecting cars and people and toasters. But manufacturing and industry account for a giant slice of global GDP. "All the appliances in your home could be wired up and monitored, but the kind of money you make in airlines or health care dwarfs that," Immelt remarked.

There is another constituency for the campaign: engineers inside GE. To them, operational software isn't anything new. Nor are control systems—even a steam locomotive has one. But here Immelt was betting they could reinvent these systems. "You do embedded systems? My God, how boring is that? It's like, put a bullet in your head," says Brian Courtney, a GE manager based in Lisle, Illinois. "Now it's the

SOURCE: GE; COURTESY OF GE

hottest job around." At the Lisle center, part of GE's Intelligent Platforms division, former field engineers sit in cubicles monitoring squiggles of data coming off turbines in Pakistan and oil rigs in onetime Soviet republics. Call this version 1.0 of the industrial Internet. On the walls, staff hang pictures of fish; each represents a problem, like a cracked turbine blade, that was caught early. More and more, GE will be using data to anticipate maintenance needs, says Courtney.

A challenge for GE is that it doesn't yet have access to most of the data its machines produce. Courtney says about five terabytes of data a day comes into GE. Facebook collects 100 times as much. According to Richard Soley, head of the Industrial Internet Consortium, a trade group GE created this year, industry has been hobbled by a "lack of Internet thinking." A jet engine has hundreds of sensors. But measurements have been collected only at takeoff, at landing, and once midflight. GE's aviation division only recently found ways to get all the flight data. "It sounds crazy, but people just didn't think about it," says Soley. "It's like the Internet revolution has just not touched the industrial revolution."

GE is trying to close that gap. Its software center in San Ramon created an adaptation of Hadoop, big-data software used by the likes of Facebook. GE also invested $100 million in Pivotal, a cloud computing company. On the crowdsourcing site Kaggle, it launched public competitions to optimize algorithms for routing airline flights, which can save fuel.

All this could sound familiar to anyone who works with consumer Internet technology, acknowledges Bernie Anger, general manager of GE's Intelligent Platforms division. But he says GE is thinking about what to do next to use connectivity, and more computers, to inject "new behavior" into machines. He gives the example of a field of wind turbines that communicate and move together in response to changes in wind. "We are moving into big data, but it's not because we want to become Google," he says. "It's because we are dramatically evolving manufacturing."

—Antonio Regalado

Emerged Technologies

The Internet of You

As wearable devices get better-looking and more powerful, we'll trust them to monitor and control more of our lives.

● The Internet of things typically conjures images of "smart" light bulbs and automatic door locks. Yet with an ever larger number of smart watches, activity trackers, and head-worn computers hitting the market, we're becoming part of the Internet of things, too.

Slowly but surely, a few wearable devices—mainly high-tech pedometers like those from Fitbit and Jawbone—are catching on with consumers, and many researchers and companies are certain that body-worn computers will become second nature—sensing, recording, and transmitting data to and from our bodies, to networks around us.

19 million

Number of watches, bands, and other wearables expected to ship in 2014

For the most part, wearables still lack wide appeal. Some, like Google Glass, elicit ambivalence. IDC estimates that manufacturers will ship 19 million watches, bands, and other wearables next year—barely a flicker next to the billion or so smartphones sold in 2013.

Wearables are still looking for their killer app. Now some people have begun to imagine that the Internet of things will provide it. If indeed our houses become filled with smart devices like door locks, a watch or wristband may be the most convenient way to control them or let them know our needs.

"Your car should know that you're tired because you didn't sleep that well, so it should be alert to that, how awake are you when you're driving, those things," says Hosain Rahman, the CEO of Jawbone, a 14-year-old company that makes earphones, speakers, as well as wrist-worn fitness trackers. "I just think that things that are on your body—wearables—ultimately will [control] all the smart stuff and be kind of at the center point."

Jawbone is among the first to try to turn a wearable into such a lifestyle remote. Jawbone's Up24 wristband can act as a trigger for the Web service IFTTT ("If This, Then That") by using its low-energy Bluetooth radio to share the data it gathers about you with an app on your smartphone. For now, it does only simple things. If you have an Internet-connected heater, the wristband can signal it to turn on when you get up in the morning. The idea is that the environment reacts to you.

Such ideas are in their infancy. Many companies are still struggling to get anyone to put a wearable computer on. Another problem is power. With Google Glass, for instance, you'll get a few hours of use before it needs to be recharged. And the biggest power draw is usually the wireless chip that lets these devices communicate. That's why MC10, a startup manufacturing soft, thin electronics, is experimenting with "every novel form of power source," says cofounder Ben Schlatka. One possibility comes from a project at Columbia University called Enhants. Researchers there are developing small, flexible tags that harvest energy from light or as they are shaken by movement. In an upcoming →

Wearable Wireless

Wireless electronic devices are gathering useful data about people.

OMSignal $240	OMShirt (gathers biometric data)
Jawbone $150	Up24 wireless wristband activity tracker
OMG Life $399	Autographer wearable camera
Google research stage	Glucose-sensing wireless contact lens
Proteus Biomedical research stage	Battery-powered drugs

SOURCE: MIT TECHNOLOGY REVIEW

research paper, they describe outfitting 40 people with flat sensors attached to different parts of their bodies while they walked, ran, or relaxed. Sensors were able to harvest enough energy to transmit data continuously at a rate of one kilobit per second. That's not much, but it could be enough for simple applications like authenticating someone's ID or reading the local temperature. Enhants researcher Peter Kinget, a professor of electrical engineering at Columbia, says enough energy can typically be harvested to wirelessly link a sensor on your body to a smartphone—something we're already comfortable carrying everywhere we go. *—Rachel Metz*

Emerged Technologies

Silicon Valley to Get a Cellular Network, Just for Things

A French company plans to build a wireless slow lane for small, low-power devices.

● San Francisco is set to get a new cellular network later this year, but it won't help fix the city's spotty mobile-phone coverage. This wireless network is exclusively for things.

The French company SigFox says it picked the Bay Area to demonstrate a wireless network intended to make it cheap and practical to link anything to the Internet, from smoke detectors to dog collars, bicycle locks, and water pipes.

Regular mobile networks are jammed with traffic from phone calls and people downloading videos. But for the Internet of things to become a reality, similar capabilities will need to be extended to billions of objects, many of them embedded in the environment and powered by small batteries. "If you want to get to billions of connections like that, you require a completely new type of network," says

Global Wireless Boom

Devices that make up the Internet of things will compete for space on networks dominated by phones, games, and video.

Wirelesss traffic, in terabytes per month

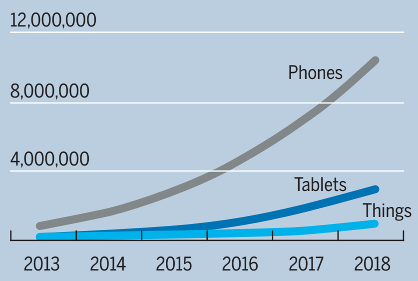

Luke D'Arcy, director of SigFox's operations in the U.S.

SigFox's network will cover the San Francisco peninsula from its urban tip to the sprawling Silicon Valley region 40 miles to the south. It will be the company's first U.S. deployment of a network technology that already covers the whole of France, most of the Netherlands, and parts of Russia and Spain. SigFox built those by adding its own equipment to existing cell towers and radio antennas. Customers include the French insurance company MAAF, which offers smoke and motion detectors that notify homeowners with a text message on their phones when a sensor is triggered or needs a new battery.

100 bits per second

Speed at which objects connected to SigFox's network can operate

The Silicon Valley network will use the unlicensed 915-megahertz spectrum band commonly used by cordless phones. Objects connected to SigFox's network can operate at very low power but will be able to transmit at only 100 bits per second—slower by a factor of 1,000 than the networks that serve smartphones. But that could be enough for many applications.

Indeed, semiconductor companies like Intel and Broadcom are also in a race to make far cheaper, far smaller, and much-lower-power wireless chips. Several showed off these "miniature computers" at the Consumer Electronics Show this year. "They saw the cell phone turn into the smartphone, and so companies are saying 'What is next?'" says David Blaauw, a professor of engineering at the University of Michigan. Blaauw builds millimeter-scale wireless computers that he believes may one day report data from just about anywhere, even from inside a patient's tumor.

A SigFox base station can serve a radius of tens of kilometers in the countryside and five kilometers in urban areas. To connect to the network, a device will need a $1 or $2 wireless chip that's compatible, and customers will pay about $1 in service charges per year per device.

By reaching into the Bay Area first (with expansion to tech hubs such as Austin, Cambridge, and Boulder in its sights), SigFox hopes to catch the interest of a region where venture capitalists poured nearly $1 billion into startup companies focusing on the Internet of things last year, according to the research firm CB Insights. One of those startups, Whistle, makes a fitness-tracking collar for dogs. It has raised $6 million and is located in a corner of San Francisco that's been called "IoT Town" thanks to its profusion of similar ventures.

Ben Jacobs, Whistle's CEO, says the collar communicates by Bluetooth to a phone, or via a home Wi-Fi router. That limits what it can do. But a new version using SigFox's technology will have a constant Internet connection anywhere in town, letting it act as a beacon for lost pets. Previously, that would have required an expensive and power-hungry cellular phone on the collar.

SigFox is in a hurry to get its network in place before competitors arrive. Jacob Sharony, a principal at the wireless consultancy Mobius Consulting, says large wireless companies are preparing machine-only networks as well, and these may operate at much higher speeds. A new long-range, low-power Wi-Fi standard that has the backing of some major U.S. companies, including Qualcomm, could hit the market in 2016. "It will likely be a major contender even though it is somewhat late to the game," says Sharony.

—Tom Simonite

SOURCE: CISCO VISUAL NETWORKING INDEX

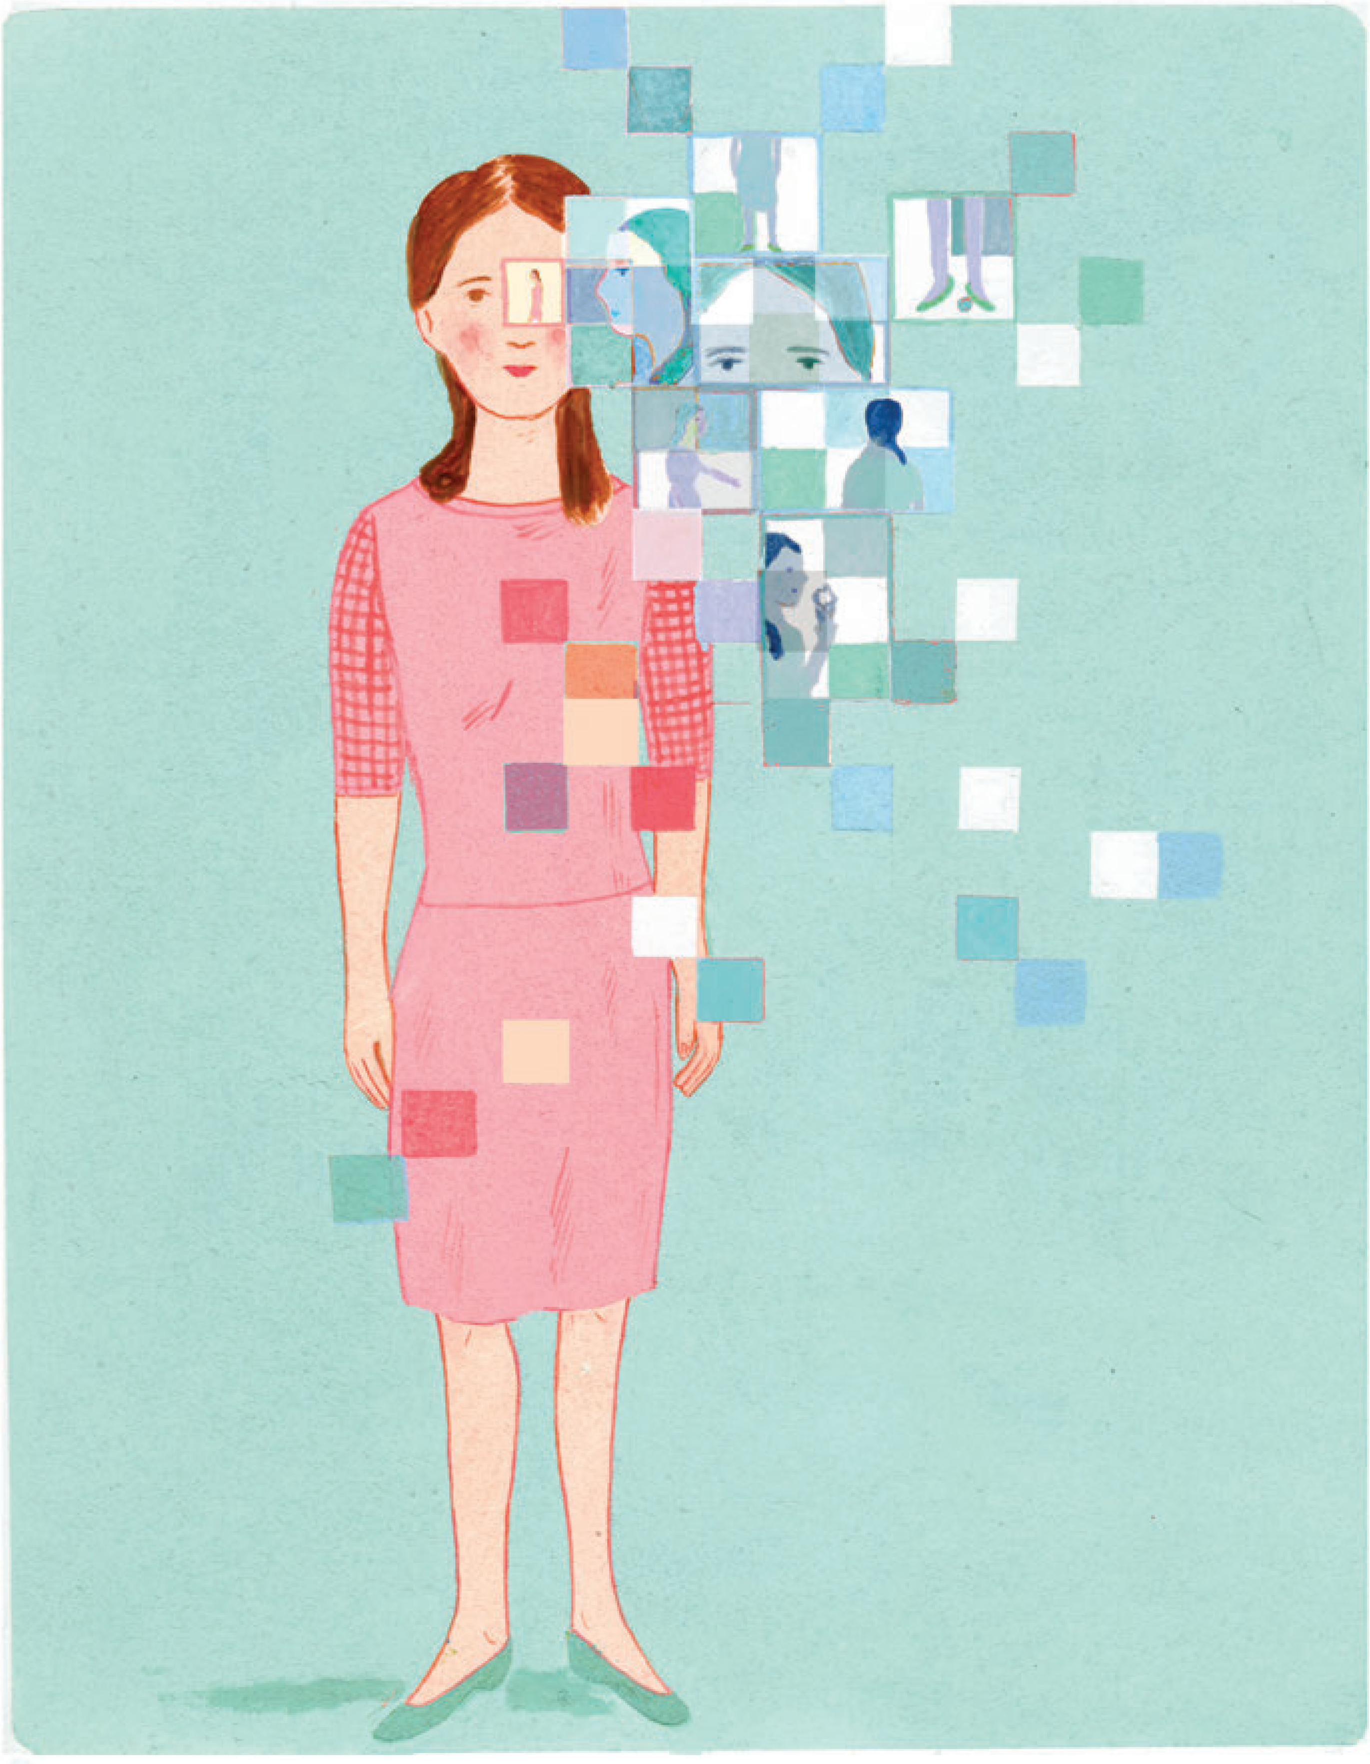

Illustration by Brian Cronin

Reviews

My Life, Logged

If a device could capture every moment in life for your easy recall later, would you want it to? There are plenty of things I'd rather forget.

By Rachel Metz

I always knew I was short, but it wasn't until recently that I realized exactly how short.

That's because I've been trying a couple of life-logging devices—gadgets that clip to my shirt or hang around my neck, automatically taking photos of the world around me. The results? Some neat shots of family, friends, and Silicon Valley life tucked in amongst countless photos of torsos, legs, sidewalks, and bicycle handlebars.

Life logging has long been an activity for a few diehard data fanatics and academics. Early adherents included University of Toronto professor Steve Mann, who as a graduate student at MIT in 1994 began wearing a wireless camera that could record images from his point of view and display them online (see "Wearable Technology as a Human Right," March/April). In 1998, Microsoft researcher Gordon Bell started collecting as much digital information about his life as he could in an effort to create a searchable archive of his memories. He even wrote a book, *Total Recall*, arguing that cataloguing everything was a better way to live and would eventually become pervasive.

We're not yet there, but we've been moving in that direction. Smartphones are widespread, with cameras and apps affording plenty of opportunities to log and share the minutiae of life, and the costs of bandwidth and storage have dropped precipitously. And now we have access to life-logging gadgets like the Autographer and the Narrative Clip: small, clip-on cameras that continually take pictures on their own. Their makers believe you want to wear a gadget that will document life for you, no effort required.

I was open to the possibility. After all, I often whip out my smartphone to capture snippets of life as it happens—firemen washing their trucks down the street from my office, flowers artfully arranged in my living room, a raccoon hanging out next to my apartment building. My aim is to share them and, occasionally, take a stroll down the digital memory lanes that are Instagram and Facebook. Perhaps a device that captures my life as it unfolds, without my having to take time out to compose images, might yield enough benefit to outweigh the annoyance of having to wear (and charge) yet another gadget. I wondered: would it make family time and bike rides more enjoyable? Would it help me remember people and events? Would I even be able to sort through the resulting deluge of photographs?

Autographer
$399

Narrative Clip
$279

Point and shoot

I hated wearing a camera at first. I felt extremely self-conscious, especially with the bulkier Autographer around my neck. It packs a wide-angle lens and sensors—including an accelerometer, magnetometer, and thermometer—into a package about the size and heft of an anchovy can. I found myself apologizing to friends and store clerks, blaming my unusual tech gear on a professional assignment, but that just called even more attention to it.

Still, it's good that the devices announce themselves. They look geeky, but not sneaky. People can (and will) ask why you have a weird square (Narrative Clip) or shiny black bar (Autographer) on your shirt collar, and especially in the case of the Autographer and its wide-angle lens, they'll be able to identify it as a camera. The Narrative Clip's lens is less obvious, but it gleams when hit by light to at least hint at what it is. If somebody has a problem with it, you can simply take it off.

And you will frequently take it off anyway. Autographer, with its sensors, GPS, and Bluetooth connection to my iPhone, sucked up battery life; I got only a few hours out of it when it was set to take photos at a "medium" frequency, which is up to 240 images per hour. The Narrative Clip was able to last more than a day, owing to the fact that it has no display and fewer sensors, and, irritatingly, plugs into a computer to sync photos.

Remembering to charge and wear the Autographer or Narrative Clip regularly felt like a chore. Though I eventually got into the habit, there was still a big problem: the content of the photos. My days often consist of hours spent gazing at a glowing monitor while my fingers tickle a keyboard. I felt embarrassed looking back at days' worth of photos that showed the same shots of my desk, computer, and office skylight.

What saved both gadgets from being tossed in a drawer was special occasions. Envisioning the ease with which they might capture time with family and friends, I took the Narrative Clip on a spring ride with a friend through the fleetingly green hillsides of Northern California and the Autographer to a playdate with my one-year-old niece.

The most interesting—and surreal—photos of the whole experiment captured my bike ride. The first few images show my garage, with bike helmets and bikes hung neatly on hooks. Then it's outside on the streets of San Francisco, followed by miles of mostly blurry trees, blue skies, and power lines—apparently, the camera was pointed nearly straight up.

In search of focus

These photos were fun to look back on, but are they really compelling enough to justify buying a camera for my lapel?

It didn't seem so, but I decided to ask someone who has devoted years to life logging. With the Narrative Clip attached to my shirt collar, I visited Gordon Bell, 79, now a researcher emeritus at Microsoft Research in San Francisco.

As part of a project called MyLifeBits, Bell archived pretty much all facets of his life except conversations (to avoid legal issues) from 1998 to 2007. He scanned photos and books, saved e-mails and instant messages, and captured the world around him by hanging a Microsoft-developed wearable camera called the SenseCam around his neck. (The Autographer, which is produced by a subsidiary of Oxford Metrics Group, a maker of imaging sensors and software, is actually based on the SenseCam technology.) MyLifeBits ended in 2007 because two of his colleagues left to join other projects,

The Narrative Clip, left, weighs 20 grams. The Autographer, which has a wide-angle lens, sensors, and display, weighs 58 grams.

COURTESY OF NARRATIVE; COURTESY OF AUTOGRAPHER

but Bell says he continues to save about as much information as he did when it was going. He had a kidney removed in April to treat cancer, and during our talk he wore a Basis health-monitoring watch that he's using to track his heart rate, which he says has risen since the operation (he suspects a beta blocker's interactions with his pacemaker).

After 15 years of life logging, he says, he often finds it helpful to consult his digital evidence store for details of transactions and "professional articles," in addition to photos of his daily whereabouts. He doesn't see the point of relying primarily on our fickle memories. "I want the real record, the real ground truth," he says. Just knowing he has these memories at hand makes him feel more confident than he was before he started tracking everything.

For the MyLifeBits project, researchers catalogued events for later recall by storing all those images and messages in a database; users could select several items and assign them keywords, or file an item in several places at once. Bell says the database was eventually discontinued because researchers didn't think many people would want to use something so complex. These days, he uses the basic desktop search on his computer. He's divided folders into "personal" and "professional" categories, with an "active" or "archival" designation within each category to make them easier to sort through. He annotates file names with metadata indicating details such as who is in a photo and what it documents.

Still, Bell acknowledges that we're far from having good ways to organize and revisit these memories. I noticed that quickly in my life-logging experiment. The iPhone apps for the Narrative Clip and the Autographer look more like photo dumping grounds than organized collections of memories. A lot of interface design and computer-vision work is needed before a life-logging device will appeal to consumers and endure over months and years.

Autographer's iPhone app hints at one potentially smart organization method: it has a map on the bottom half of the display, and as you browse images, it shows a dot indicating where each photo was taken. Tapping a picture calls up detailed sensor information gathered at the time the image was snapped—including the brightness of the light, the temperature, and the direction of the camera. A lot of this data is just noise, but seeing the geographic location of photos was valuable because it jogged my memory: "Oh yeah, that's where I was when that happened."

The value of forgetting

No matter how organized these tools may become, though, there's something I just can't get past: sometimes forgetting is even more helpful than remembering.

Anind Dey, an associate professor of human-computer interaction at Carnegie Mellon University, has been studying how self-monitoring or life logging can serve as a memory aid for Alzheimer's patients or a tool to learn more about what's going on in the lives of people with autism. Beyond that, he says, "I still think this is a niche market."

One big reason is that retaining everything we experience "doesn't really match the way we think," Dey says. "I think of bad incidents in my life, and in my head I've made them better, or I'd be really unhappy."

Me too. While I'm happy recalling events like the bike ride with my friend, there were some moments captured by my life-logging cameras that I'd rather not relive, like one particularly stressful Saturday night I spent at the wood shop helping my fiancé, Noah, finish up a project. The Narrative Clip's camera captured hours of exhaustion and irritation as we fumbled to glue little pieces of wood into small slots cut out of a giant map. I had pretty much forgotten about that miserable night until I glanced back at a set of images that show Noah in the corner of the frame with a sad look on his face. Every time I look at it, I wince. And this wouldn't even register on my lifetime scale of awful memories. It's hard to imagine how much more difficult it would be to face photo evidence of, say, a friend's or family member's death.

The Narrative Clip lets you swipe to the left on any collection of images to reveal a delete button; Autographer makes it more complicated. Either way, I have to make a decision to delete these images, which brings the unpleasant memories to mind again.

I asked Bell whether recording everything makes him more likely to remember the things he'd rather forget. Laughing, he suggested that rather than deleting something, I put it in a folder labeled "Don't Ever Look at This, and Forget It." He's not joking. Although he realizes there will be some photographs you never want to see again, he still thinks these logged memories can be useful: "The value, I think, is—well, this is something you can give to your therapist."

I'd rather continue with my targeted approach to capturing, organizing, and sharing my life. There's something nice and deliberate about taking a photo, even if I'm just pressing a virtual shutter button on my iPhone screen. Gadgets like the Autographer and the Narrative Clip take some of the work out of it, but like me, they come up short.

Rachel Metz is IT editor for Web and social media at MIT Technology Review.

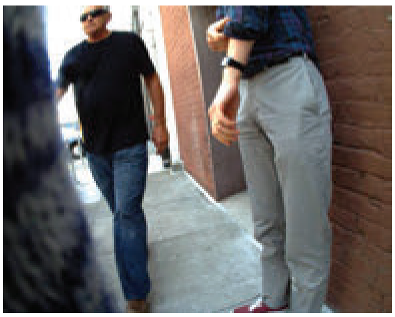

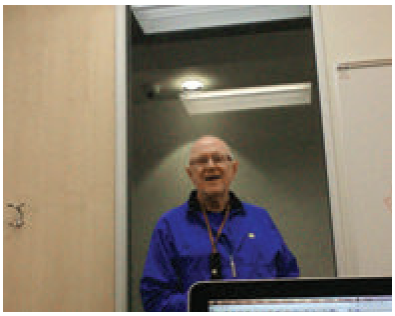

RACHEL METZ/MIT TECHNOLOGY REVIEW

Clockwise from top left: an awkward shot taken by the Autographer; the Autographer's view from a bike; a Narrative Clip picture of Gordon Bell.

JAN BAJTLIK

A Plan B for Climate Agreements

U.N. negotiations are going nowhere, and greenhouse-gas emissions are soaring. It's time to move on.

By Kevin Bullis

In 2007, just before he accepted the Nobel Prize on behalf of the U.N.'s Intergovernmental Panel on Climate Change (IPCC), Rajendra Pachauri, the organization's leader, declared that the world was running out of time to prevent catastrophic global warming. "If there's no action before 2012, that's too late," Pachauri told the *New York Times*. "What we do in the next two or three years will determine our future. This is the defining moment."

This April, the IPCC released a long-awaited report assessing just how far we've come since Pachauri's stark pronouncement. The news was grim. There has yet to be any sign of the global action that Pachauri and others had desperately sought. In 2007, the IPCC called for emissions to level off by 2015, but the world is emitting greenhouse gases faster than ever. Even now, Pachauri and some other IPCC leaders remain publicly optimistic, saying it's still possible to avoid catastrophic climate change if we act "very soon." But delve into the new IPCC report itself and you'll find a much less hopeful picture.

"What the report overall shows is that the only way you're going to stop climate change is by assuming that governments will make a whole series of heroic efforts," says David Victor, director of the Laboratory on International Law and Regulation at the University of California at San Diego and one of the lead authors of the report. For example, if we're to meet the U.N. goal of limiting the increase in the world's average temperature to 2 °C above preindustrial levels, all key technologies—wind, solar, nuclear, power plants that capture and store carbon dioxide, and so on—have to scale up quickly, even though some haven't been growing at all and carbon capture at power plants has yet to be commercially deployed. The report says the world's governments must also immediately agree to binding and effective climate policies—even though massive efforts over the past two decades have failed to produce any such agreements.

Climate Change 2014: Mitigation of Climate Change
Working Group III Contribution to the Intergovernmental Panel on Climate Change's Fifth Assessment Report

"A lot of publicity was given to the fact that the IPCC says limiting warming is doable," says Robert Stavins, director of the Harvard Environmental Economics Program and also a key IPCC author. But the situation "is not as rosy as it's been characterized," he says. "When you introduce political reality, then you're talking about [2 °C] not being feasible."

Achieving consensus among nearly 200 countries is one challenge for climate policy, but it's not necessarily the main one. Another problem is the U.N.'s focus on emissions limits.

The report inadvertently suggests another conclusion: after more than two decades of the U.N. climate-treaty process, long thought to be our best shot at getting governments to act on global warming, it's time for a new approach. Although there are many reasons emissions have continued to increase, an obvious one is simply that the U.N. approach to climate change—which involves gathering representatives from nearly 200 countries and trying to hash out treaties that articulate global, binding limits on greenhouse gases—isn't working. The new IPCC report itself notes that governments are increasingly turning to forums outside the U.N. to make progress.

Out of control

The U.N. became seriously involved in addressing climate change in 1988, when it created the IPCC, an organization tasked with assessing the science on climate change and laying the foundation for climate treaties. The IPCC produces voluminous reports every six years or so; the report in April was part of the organization's fifth such assessment. An organization designed to actually produce climate policy came later, in 1994, when the U.N. established the Framework Convention on Climate Change, which laid ground rules for international coöperation on climate policy and set the vague yet ambitious goal of stabilizing greenhouse gases "at a level that would prevent dangerous anthropogenic (human-induced) interference with the climate system."

The most famous agreement to emerge so far is the Kyoto Protocol, a 1997 treaty that established binding emissions limits on some industrial countries and set up a mechanism to finance emissions-reducing projects in poor countries. Kyoto expired in 2012 and has yet to be replaced. The new IPCC report concludes the obvious, stating that Kyoto accomplished very little—the economic collapse in Eastern Europe after the fall of the Soviet Union did more to reduce emissions, according to Victor. Since Kyoto, U.N. actions have been largely symbolic; an example is the nonbinding emissions targets that came out of a meeting in Copenhagen in 2009. The goal of a global, all-inclusive treaty keeps getting kicked down the road. The most recent meeting, in Warsaw, ended with nations promising they'd come up with a universal agreement in Paris next year. Don't hold your breath.

Achieving consensus among nearly 200 countries is one challenge for the United Nations on climate policy, but it's not necessarily the main one. Another problem is the U.N.'s focus thus far on specific emissions limits. Governments don't know how much it will cost to comply with these limits, because it's often unclear how much low-carbon technologies will cost to deploy. Climate negotiators don't want to commit their governments to a treaty whose economic effects are unpredictable.

What's more, not every government is capable of enforcing such regulations. Even countries with significant regulatory powers would struggle to monitor and control their overall carbon dioxide emissions. As an example, the United States recently set aggressive fuel-economy standards for cars. These could reduce emissions, but total emissions ultimately depend on how much people drive.

A far more effective strategy for reducing carbon dioxide emissions would be to encourage international coöperation on investing in new, cleaner energy technologies and to establish incentives for such investments. But instead, the U.N. has obsessed over setting specific binding emissions limits on greenhouse gases.

Changing climate

There is no simple alternative to the U.N. process. But to stand a chance of being adopted and making an impact, international climate policy will probably have to start with smaller groups of countries and focus on single industries or economic sectors. These kinds of policies won't be enough to stabilize greenhouse-gas levels on their own, but they could slow global warming in the near term. And collaborations between countries could lay a foundation for more ambitious policies when the costs of low-carbon technologies come down.

The U.S. and China, which can afford to spend money on R&D and demonstration projects for new energy technologies, need to redouble their efforts. One important area for investment is technology for capturing carbon dioxide from fossil-fuel power plants and permanently storing it. Carbon capture and storage hasn't been demonstrated anywhere at a large scale, yet the IPCC says that without it, the cost of stabilizing greenhouse-gas levels could double. At least for some key technologies, R&D collaborations are happening now: the U.S. and China are collaborating on advanced nuclear power plants that could be cheaper and safer. That kind of coöperation could serve as a model for future projects.

Even a handful of countries, such as the U.S., China, India, and members of the European Union, could make a big impact. Ten countries account for nearly 70 percent of the world's total emissions. If China, the world's largest emitter, were to take measures to reduce emissions (and it's starting to, in part to cut pollution), it could then plausibly influence other countries—particularly its key trading partners—to implement technology changes.

Meanwhile, large, wealthy countries could help poor countries adopt greener policies as a way to reap immediate health benefits. For example, the IPCC report notes work by a consortium of countries to find ways to promote the use of cleaner cookstoves in poor countries. Conventional stoves used in such countries emit soot, which causes respiratory problems and warms the atmosphere by absorbing sunlight. Reducing soot emissions would have an immediate impact on warming, because soot—unlike carbon dioxide, which lingers in the atmosphere for centuries—disappears quickly after emissions stop.

Actions like these don't have to replace the U.N. process—they can easily happen alongside it. But diplomats might want to spend less time on U.N. talks that aren't making progress and more on smaller efforts that can. Going forward, the U.N. should focus on things it does well. It may not be a good forum for hashing out universal agreements that affect every economy in the world, but as countries forge new agreements, the U.N. can verify that they're meeting their obligations.

Simply investing in technology and establishing piecemeal policies won't limit warming to two degrees. The IPCC report suggests it's already too late for that. But unlike the U.N.'s approach, these efforts could at least make tangible progress. The U.N.'s approach isn't working. It's time to recognize that and move on.

Kevin Bullis is MIT Technology Review*'s senior editor for energy.*

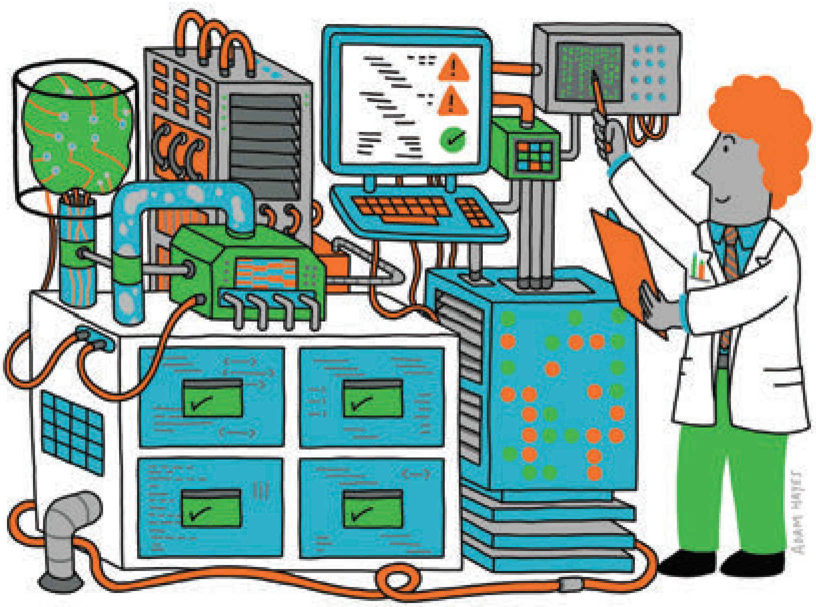

Imposing Security

Computer programmers won't stop making dangerous errors on their own. It's time they adopted an idea that makes the physical world safer.

By Simson L. Garfinkel

THREE COMPUTER BUGS THIS YEAR exposed passwords, e-mails, financial data, and other kinds of sensitive information connected to potentially billions of people. The flaws cropped up in different places—the software running on Web servers, iPhones, the Windows operating system—but they all had the same root cause: careless mistakes by programmers.

Each of these bugs—the "Heartbleed" bug in a program called OpenSSL, the "goto fail" bug in Apple's operating systems, and a so-called "zero-day exploit" discovered in Microsoft's Internet Explorer—was created years ago by programmers writing in C, a language known for its power, its expressiveness, and the ease with which it leads programmers to make all manner of errors. Using C to write critical Internet software is like using a spring-loaded razor to open boxes—it's really cool until you slice your fingers.

Alas, as dangerous as it is, we won't eliminate C anytime soon—programs written in C and the related language C++ make up a large portion of the software that powers the Internet. New projects are being started in these languages all the time by programmers who think they need C's speed and think they're good enough to avoid C's traps and pitfalls.

"Heartbleed"
flaw in OpenSSL

"Goto fail"
flaw in Apple operating systems

CVE-2014-1776
flaw in Internet Explorer

But even if we can't get rid of that language, we can force those who use it to do a better job. We would borrow a concept used every day in the physical world.

Obvious in retrospect

Of the three flaws, Heartbleed was by far the most significant. It is a bug in a program that implements a protocol called Secure Sockets Layer/Transport Layer Security (SSL/TLS), which is the fundamental encryption method used to protect the vast majority of the financial, medical, and personal information sent over the Internet. The original SSL protocol made Internet commerce possible back in the 1990s. OpenSSL is an open-source implementation of SSL/TLS that's been around nearly as long. The program has steadily grown and been extended over the years.

Today's cryptographic protocols are thought to be so strong that there is, in practice, no way to break them. But Heartbleed made SSL's encryption irrelevant. Using Heartbleed, an attacker anywhere on the Internet could reach into the heart of a Web server's memory and rip out a little piece of private data. The name doesn't come from this metaphor but from the fact that Heartbleed is a flaw in the "heartbeat" protocol Web browsers can use to tell Web servers that they are still connected. Essentially, the attacker could ping Web servers in a way that not only confirmed the connection but also got them to spill some of their contents. It's like being able to check into a hotel that occasionally forgets to empty its rooms' trash cans between guests. Sometimes these contain highly valuable information.

Heartbleed resulted from a combination of factors, including a mistake made by a volunteer working on the OpenSSL program when he implemented the heartbeat protocol.

NEIL WEBB


A PLACE OF
INSPIRATION

Although any of the mistakes could have happened if OpenSSL had been written in a modern programming language like Java or C♯, they were more likely to happen because OpenSSL was written in C.

Apple's flaw came about because some programmer inadvertently duplicated a line of code that, appropriately, read "goto fail." The result was that under some conditions, iPhones and Macs would silently ignore errors that might occur when trying to ascertain the legitimacy of a website. With knowledge of this bug, an attacker could set up a wireless access point that might intercept Internet communications between iPhone users and their banks, silently steal usernames and passwords, and then reëncrypt the communications and send them on their merry way. This is called a "man-in-the-middle" attack, and it's the very sort of thing that SSL/TLS was designed to prevent.

Remarkably, "goto fail" happened because of a feature in the C programming language that was known to be problematic *before C was even invented!* The "goto" statement makes a computer program jump from one place to another. Although such statements are common inside the computer's machine code, computer scientists have tried for more than 40 years to avoid using "goto" statements in programs that they write in so-called "high-level language." Java (designed in the early 1990s) doesn't have a "goto" statement, but C (designed in the early 1970s) does. Although the Apple programmer responsible for the "goto fail" problem could have made a similar mistake without using the "goto" statement, it would have been much less probable.

We know less about the third bug because the underlying source code, part of Microsoft's Internet Explorer, hasn't been released. What we do know is that it was a "use after free" error: the program tells the operating system that it is finished using a piece of memory, and then it goes ahead and uses that memory again. According to the security firm FireEye, which tracked down the bug after hackers started using it against high-value targets, the flaw had been in Internet Explorer since August 2001 and affected more than half of those who got on the Web through traditional PCs. The bug was so significant that the Department of Homeland Security took the unusual step of telling people to temporarily stop running Internet Explorer. (Microsoft released a patch for the bug on May 1.)

Automated inspectors

There will always be problems in anything designed or built by humans, of course. That's why we have policies in the physical world to minimize the chance for errors to occur and procedures designed to catch the mistakes that slip through.

Home builders must follow building codes, which regulate which construction materials can be used and govern certain aspects of the building's layout—for example, hallways must reach a minimum width, and fire exits are required. Building inspectors visit the site throughout construction to review the work and make sure that it meets the codes. Inspectors will make contractors open up walls if they've installed them before getting the work inside inspected.

The world of software development is completely different. It's common for developers to choose the language they write in and the tools they use. Many developers design their own reliability tests and then run the tests themselves! Big companies can afford separate quality-assurance teams, but many small firms go without. Even in large companies, code that seems to work properly is frequently not tested for lurking security flaws, because manual testing by other humans is incredibly expensive—sometimes more expensive than writing the original software, given that testing can

reveal problems the developers then have to fix. Such flaws are sometimes called "technical debt," since they are engineering costs borrowed against the future in the interest of shipping code now.

The solution is to establish software building codes and enforce those codes with an army of unpaid inspectors.

Crucially, those unpaid inspectors should not be people, or at least not only people. Some advocates of open-source software subscribe to the "many eyes" theory of software development: that if a piece of code is looked at by enough people, the security vulnerabilities will be found. Unfortunately, Heartbleed shows the fallacy in this argument: though OpenSSL is one of the most widely used open-source security programs, it took paid security engineers at Google and the Finnish IT security firm Codenomicon to find the bug—and they didn't find it until two years after many eyes on the Internet first got access to the code.

Many developers design their own reliability tests and then run the tests themselves. Even in large companies, code that seems to work properly is frequently not tested for lurking flaws.

Instead, this army of software building inspectors should be software development tools—the programs that developers use to create programs. These tools can needle, prod, and cajole programmers to do the right thing.

This has happened before. For example, back in 1988 the primary infection vector for the world's first Internet worm was another program written in C. It used a function called "gets()" that was common at the time but is inherently insecure. After the worm was unleashed, the engineers who maintained the core libraries of the Unix operating system (which is now used by Linux and Mac OS) modified the gets() function to make it print the message "Warning: this program uses gets(), which is unsafe." Soon afterward, developers everywhere removed gets() from their programs.

The same sort of approach can be used to prevent future bugs. Today many software development tools can analyze programs and warn of stylistic sloppiness (such as the use of a "goto" statement), memory bugs (such as the "use after free" flaw), or code that doesn't follow established good-programming standards. Often, though, such warnings are disabled by default because many of them can be merely annoying: they require that code be rewritten and cleaned up with no corresponding improvement in security. Other bug-finding tools aren't even included in standard development tool sets but must instead be separately downloaded, installed, and run. As a result, many developers don't even know about them, let alone use them.

To make the Internet safer, the most stringent checking will need to be enabled by default. This will cause programmers to write better code from the beginning. And because program analysis tools work better with modern languages like C# and Java and less well with programs written in C, programmers should avoid starting new projects in C or C++—just as it is unwise to start construction projects using old-fashioned building materials and techniques.

Programmers are only human, and everybody makes mistakes. Software companies need to accept this fact and make bugs easier to prevent.

Simson L. Garfinkel is a contributing editor to MIT Technology Review *and a professor of computer science at the Naval Postgraduate School.*

53 Years Ago

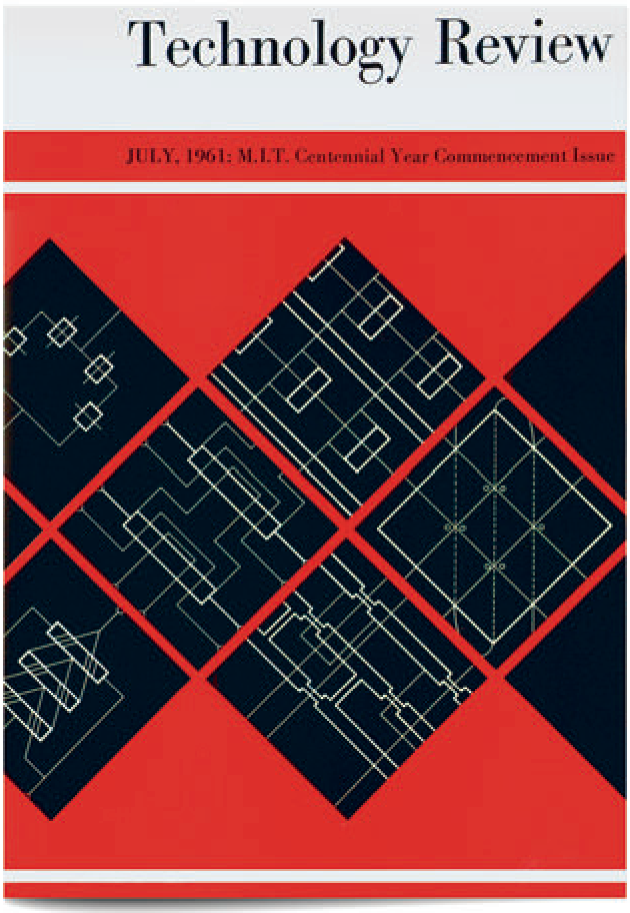

Reading the Mind

A 1961 essay speculated on where research into the physical basis of thinking and communication might eventually lead.

"Biologists of my generation have dealt effectively with two major problems: 1) the search for physical and mental health and the conquest of disease and aging; 2) the search for an understanding of the biochemical foundations of life. We are now beginning a third: the search for the physical basis of mind.

To communicate with his fellows, man now transduces his thoughts to spoken or written symbols. These are reasonably satisfactory for simple messages, but inadequate for conveying complex conceptual ideas, human emotion, and spirit. Will biophysical research on mind pave the way for bypassing sensory mechanisms? It may not be unreasonable to imagine that this might eventually occur, perhaps at first requiring instrumental prosthetic aids. Pooling the diversity of individuals' learning and endowments by such interpersonal communication could inaugurate a new hierarchy of intelligence and a new kind of science. Other implications of human interthinking as a new advance in evolution have been projected by Teilhard de Chardin and by other speculative thinkers.

All advances in our understanding of mental processes, as of other natural phenomena, are made through science, and therefore do not directly touch the ontological problem of the nature of the inner self. Implied here is no attempt by research or sheer intellectual genius to grasp reality by its quantized forelock, no suggestion that man's mind is no more than a quantum mechanical automaton. On the contrary, even such revolutionary discoveries as are here projected would still be science, therefore susceptible, like all scientific endeavor, to beneficent application—but also to ultimate desecration!"

Excerpted from "Life, Science, and Inner Commitment," by Francis O. Schmitt, in the July 1961 issue of Technology Review.

MIT Technology Review (ISSN 1099-274X), July/August 2014 Issue, Reg. U.S. Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2014. 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 Dept., PO Box 16327, North Hollywood, CA 91615, or via the Internet at www.technologyreview.com/customerservice. Basic subscription rates: $29.95 per year within the United States; in all other countries, US$39.95. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051 Fort Erie, ON L2A 6C7. Printed in U.S.A. Audited by the Alliance for Audited Media