IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 06.17

LEARNING TO LOVE THINKING MACHINES
Why the brain will drive computing's next era
P. 20

HOW TO MAKE AN ARTIFICIAL BRAIN
A plan to make it happen in 30 years
P. 44

IBM BETS ON BRAIN-INSPIRED CHIPS
Soon we'll know if all the hype was justified
P. 50

CAN A MACHINE BE CONSCIOUS?
Yes, but not if it's built like today's computers
P. 62

Can We Copy the Brain?

Intensive efforts to understand and re-create human cognition will spark a new age of machine intelligence and transform the way we work, learn, and play

A SPECIAL REPORT

IEEE

CLICK HERE TO ACCESS THE JUNE ISSUE OF the institute

IEEE SPECTRUM

SPECIAL REPORT

CAN WE COPY THE BRAIN?

1 A UNIQUE MACHINE

2 THE MECHANICS OF THE MIND

3 ENGINEERING COGNITION

On the cover
Photograph for *IEEE Spectrum* by Dan Saelinger

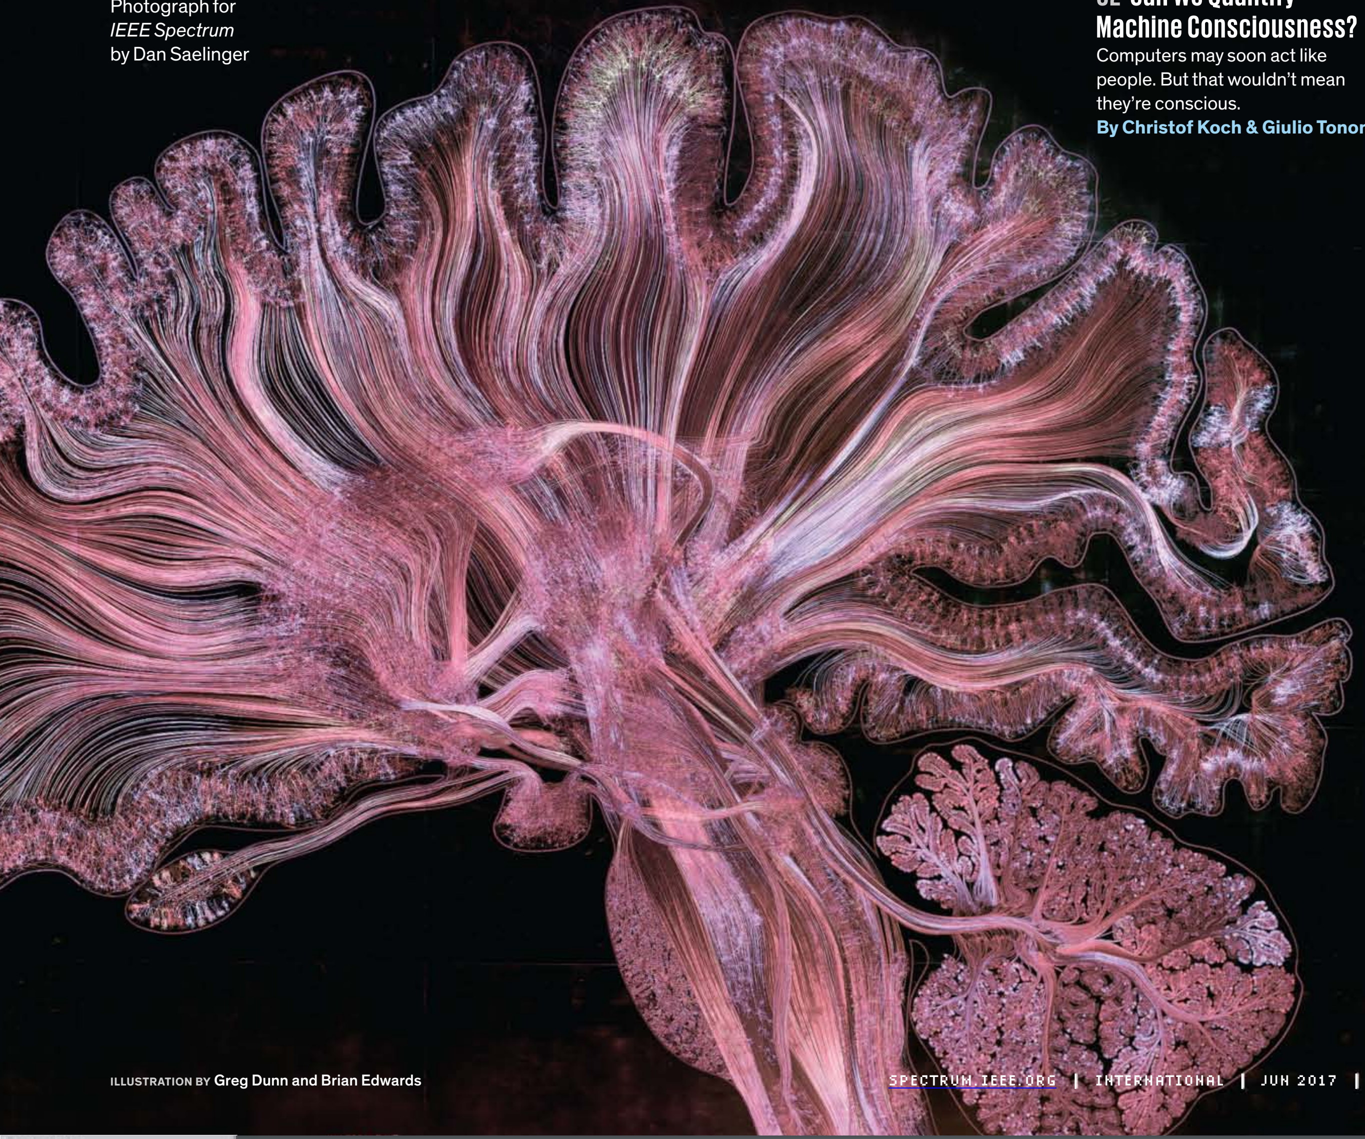

ILLUSTRATION BY Greg Dunn and Brian Edwards


AVNET®
Reach Further™

DEPARTMENTS_06.17

IEEE
SPECTRUM

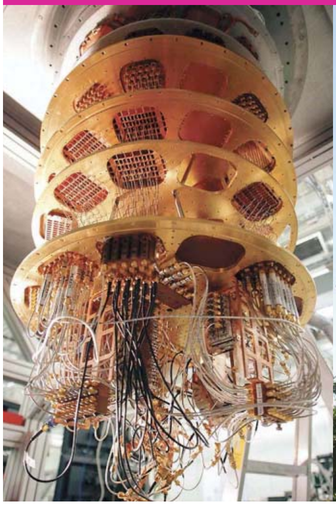

FROM LEFT: ERIK LUCERO; HENRY SCHNEIDER; STEVEN J. KEATING/SCIENCE ROBOTICS

Online
Spectrum.ieee.org
Buildings on Demand
A construction robot uses 3D printing to build large structures safely, quickly, and autonomously. It can embed rebar, fuse metal chains, and even create sculptures out of ice. Watch it 3D-print a building out of spray foam! http://spectrum.ieee.org/3dprintedbuilding0617

ADDITIONAL RESOURCES

Tech Insider / Webinars
Available at spectrum.ieee.org/webinars

- Simulation-Driven Motor Design for Vehicle Electrification—15 June

White Papers
Available at http://spectrum.ieee.org/whitepapers

- DC Motor Selection for Dynamic Motion Control Applications
- Control Design Guide for Next-Generation Machines
- Multi-Tone: Testing, Theory, and Practice Application Note
- Learn the Basics of Power Amplifier and Front-End Module Measurements
- Accelerating Aircraft Systems Development With Model-Based Engineering
- Getting Started With a Part-Centric BOM

The Institute
Available at theinstitute.ieee.org

AN ACCESSIBLE WORLD In this issue, we feature several IEEE members who are developing technologies to help people with disabilities. A senior member is working on robots to help paralyzed people live independently, while an IEEE Fellow is developing a robotic wheelchair that can tackle uneven terrain. We also feature student members who are retrofitting toy cars for children who are visually impaired and a startup that makes 3D tactile artworks for the blind.

MAKING STRIDES Member Conor Walsh has built a soft, lightweight exosuit that helps stroke patients and others with mobility issues walk with less difficulty.

MEET THE CANDIDATES The 2018 IEEE president-elect candidates, Fellow Vincenzo Piuri and Life Fellow Jacek M. Zurada, discuss their priorities for the organization.

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2017 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 54, issue No. 6, International edition. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

BACK STORY_


ALLEN INSTITUTE

A Boat of Glass and Steel

HOW MANY TYPES OF NERVE CELLS are in the human brain? "We don't know," says Christof Koch. "Probably thousands."

We'll have an answer soon, thanks in part to Koch, who with Giulio Tononi wrote "Can We Quantify Machine Consciousness?," in this issue. As president and chief scientific officer of the Allen Institute for Brain Science, in Seattle, Koch oversees a staff of more then 300 people, some of whom are now producing a complete catalog of the many kinds of neurons in the brain.

Koch [above], well known for his work in the field of consciousness, took the helm of the Brain Science institute after spending nearly three decades as a professor of both biology and engineering at the California Institute of Technology. So he was very familiar with how basic research gets done at universities.

The Allen Institute is different, though: "We're somewhere between biotech and academia," says Koch, who explains that his institute tackles problems that might require a hundred or more researchers working for several years–much beyond what a university lab could do.

At the end of 2015, Koch and his colleagues moved into a newly built 25,000-square-meter facility, donated by Paul Allen. The modern glass building was specially designed to foster collaboration and enhance esprit de corps within the institute's large research teams.

To highlight the importance of team building, Koch prefaced a recent scholarly article with a quotation from Daniel James Brown's *The Boys in the Boat*: "The team effort–the perfectly synchronized flow of muscle, oars, boat and water; the single, whole, unified and beautiful symphony that a crew in motion becomes–is all that matters." The same is true, according to Koch, when your boat is made of glass and steel and your goal is scientific understanding. ■

CITING ARTICLES IN IEEE SPECTRUM *IEEE Spectrum* publishes an international and a North American edition, as indicated at the bottom of each page. Both have the same editorial content, but because of differences in advertising, page numbers may differ. In citations, you should include the issue designation. For example, Past Forward is in *IEEE Spectrum*, Vol. 54, no. 6 (INT), June 2017, p. 72, or in *IEEE Spectrum*, Vol. 54, no. 6 (NA), June 2017, p. 76.

IEEE
SPECTRUM

EDITOR IN CHIEF
Susan Hassler, s.hassler@ieee.org
EXECUTIVE EDITOR
Glenn Zorpette, g.zorpette@ieee.org
EDITORIAL DIRECTOR, DIGITAL
Harry Goldstein, h.goldstein@ieee.org
MANAGING EDITOR
Elizabeth A. Bretz, e.bretz@ieee.org
SENIOR ART DIRECTOR
Mark Montgomery, m.montgomery@ieee.org
SENIOR EDITORS
Stephen Cass (Resources), cass.s@ieee.org
Erico Guizzo (Digital), e.guizzo@ieee.org
Jean Kumagai, j.kumagai@ieee.org
Samuel K. Moore (News), s.k.moore@ieee.org
Tekla S. Perry, t.perry@ieee.org
Philip E. Ross, p.ross@ieee.org
David Schneider, d.a.schneider@ieee.org
DEPUTY ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSOCIATE ART DIRECTOR Erik Vrielink, e.vrielink@ieee.org
SENIOR ASSOCIATE EDITORS
Rachel Courtland, r.courtland@ieee.org
Eliza Strickland, e.strickland@ieee.org
ASSOCIATE EDITORS
Celia Gorman (Multimedia), celia.gorman@ieee.org
Willie D. Jones (News), w.jones@ieee.org
Amy Nordrum, a.nordrum@ieee.org
SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
EDITORIAL RESEARCHER Alan Gardner, a.gardner@ieee.org
ADMINISTRATIVE ASSISTANT
Ramona L. Foster, r.foster@ieee.org
VIDEO INTERN Alyssa Pagano
CONTRIBUTING EDITORS Evan Ackerman, Mark Anderson, John Blau, Robert N. Charette, Peter Fairley, Tam Harbert, Mark Harris, David Kushner, Robert W. Lucky, Paul McFedries, Prachi Patel, Richard Stevenson, Lawrence Ulrich, Paul Wallich

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
EDITORIAL & WEB PRODUCTION MANAGER Roy Carubia
SENIOR ELECTRONIC LAYOUT SPECIALIST Bonnie Nani
PRODUCT MANAGER, DIGITAL Shannan Brown
WEB PRODUCTION COORDINATOR Jacqueline L. Parker
MULTIMEDIA PRODUCTION SPECIALIST Michael Spector
ADVERTISING PRODUCTION +1 732 562 6334
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org
SENIOR ADVERTISING PRODUCTION COORDINATOR
Nicole Evans Gyimah, n.gyimah@ieee.org

EDITORIAL ADVISORY BOARD
Susan Hassler, *Chair*; David Brock, Sudhir Dixit, Limor Fried, Robert Hebner, Joseph J. Helble, Grant Jacoby, Leah Jamieson, Jelena Kovacevic, Deepa Kundur, Norberto Lerendegui, Steve Mann, Allison Marsh, Jacob Østergaard, Umit Ozguner, Thrasos Pappas, H. Vincent Poor, John Rogers, Jonathan Rothberg, Umar Saif, Takao Someya, Maurizio Vecchione, Yu Zheng, Kun Zhou, Edward Zyszkowski

MANAGING DIRECTOR, PUBLICATIONS Michael B. Forster

EDITORIAL CORRESPONDENCE
IEEE Spectrum, 3 Park Ave., 17th Floor,
New York, NY 10016-5997
TEL: +1 212 419 7555 FAX: +1 212 419 7570
BUREAU Palo Alto, Calif.; Tekla S. Perry **+1 650 752 6661**

DIRECTOR, BUSINESS DEVELOPMENT,
MEDIA & ADVERTISING Mark David

ADVERTISING INQUIRIES
IEEE GLOBALSPEC
30 Tech Valley Dr., Suite 102, East Greenbush, NY 12061
+1 844 300 3098 Toll-free: +1 800 261 2052
www.globalspec.com
VP, DIGITAL MEDIA & ENGINEERING INSIGHT Don Lesem
+1 518 238 6514, don.lesem@ieeeglobalspec.com
VP, SALES & CUSTOMER CARE Peter Hauhuth
+1 303 594 8007, peter.hauhuth@ieeeglobalspec.com
SENIOR DIRECTOR, PRODUCT MANAGEMENT
& MARKETING Christian Noe
+1 518 238 6611, christian.noe@ieeeglobalspec.com
SENIOR PRODUCT MANAGER Linda Uslaner
+1 518 238 6527, linda.uslaner@ieeeglobalspec.com

REPRINT SALES +1 212 221 9595, ext. 319
REPRINT PERMISSION / LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Managing Editor, *IEEE Spectrum*.

CONTRIBUTORS_

Lee Gomes

To report "The Neuromorphic Chip's Make-or-Break Moment" [p. 50], freelance journalist Gomes set out for the flagship conference on neuromorphic computing. There, he spoke with 20 experts about where the field is headed. Gomes, who has grown leery of high-flying claims after three decades as a technology reporter, found this group to be refreshingly grounded. "I was surprised at how forthright they were about what they still had to prove to the world," he says.

Jennifer Hasler

A professor of electrical and computer engineering at the Georgia Institute of Technology, Hasler works on "physical"–including analog and neurobiologically inspired–computing, which she writes about in "A Road Map for the Artificial Brain" [p. 44]. Hasler, who trained with neuromorphic computing pioneer Carver Mead, recently developed field-programmable analog arrays, an analog and mixed-signal answer to FPGAs that can consume less than a milliwatt of power.

Jeff Hawkins

Hawkins was the founder or cofounder of Palm Computing, Handspring, and the Redwood Center for Theoretical Neuroscience. His latest company, Numenta, produces among other things an open-source computing platform based on the principles of the human neocortex. With science writer Sandra Blakeslee, he wrote *On Intelligence*. In this issue, he describes the three features of the human brain that machines will need to emulate to be considered truly intelligent [p. 32].

Karlheinz Meier

Meier is a professor of experimental physics at Heidelberg University, in Germany, who worked for more than 30 years in experimental particle physics before shifting his focus in 2005 to neuromorphic computing. He is a codirector in the European Union's Human Brain Project. In "The Brain as Computer" [p. 26], Meier writes about the many advantages to be gained by pursuing brain-inspired computing.

Fred Rothganger

Rothganger is a senior member of the technical staff at Sandia National Laboratories, in Albuquerque, N.M. He works in the labs' neural computing group, developing tools to model the brain on supercomputers. In his spare time, he writes science fiction. His novel *SuSAn: A Romance With Technology*, is about the creation of the world's first sentient robot, and he explores the real-world implications of such thinking machines in this issue [p. 20].

Can We Copy the Brain?

Massive efforts to better understand the human brain will deliver on the original promise of computer science: machines that think like us

In the mid-1940s, a few brilliant people drew up the basic blueprints of the computer age. They conceived a general-purpose machine based on a processing unit made up of specialized subunits and registers, which operated on stored instructions and data. Later inventions–transistors, integrated circuits, solid-state memory–would supercharge this concept into the greatest tool ever created by humankind.

So here we are, with machines that can churn through tens of quadrillions of operations per second. We have voice-recognition-enabled assistants in our phones and homes. Computers routinely thrash us in our ancient games. And yet we still don't have what we want: machines that can communicate easily with us, understand and anticipate our needs deeply and unerringly, and reliably navigate our world.

Now, as Moore's Law seems to be starting some sort of long goodbye, a couple of themes are dominating discussions of computing's future. One centers on quantum computers and stupendous feats of decryption, genome analysis, and drug development. The other, more interesting vision is of machines that have something like human cognition. They will be our intellectual partners in solving some of the great medical, technical, and scientific problems confronting humanity. And their thinking may share some of the fantastic and maddening beauty, unpredictability, irrationality, intuition, obsessiveness, and creative ferment of our own.

In this issue, we consider the advent of neuromorphic computing and its prospects for ushering in a new age of truly intelligent machines. It is already a sprawling enterprise, being propelled in part by massive research initiatives in the United States and Europe aimed at plumbing the workings of the human brain. Parallel engineering efforts are now applying some of that knowledge to the creation of software and specialized hardware that "learn"–that is, get more adept–by repeated exposure to computational challenges.

Brute speed and clever algorithms have already produced machines capable of equaling or besting us at activities we've long thought of as deeply human: not just poker and Go but also stock picking, language translation, facial recognition, drug discovery and design, and the diagnosis of several specific diseases. Pretty soon, speech recognition, driving, and flying will be on that list, too.

The emergence of special-purpose hardware, such as IBM's TrueNorth chips and the University of Manchester's SpiNNaker, will eventually make the list longer. And yet, our intuition (which for now remains uniquely ours) tells us that even then we'll be no closer to machines that can, through learning, become capable of making their way in our world in an engaging and yet largely independent way.

To produce such a machine we will have to give it common sense. If you act erratically, for example, this machine will recall that you're going through a divorce and subtly change the way it deals with you. If it's trying to deliver a package and gets no answer at your door, but hears a small engine whining in your backyard, it will come around to see if there's a person (or machine) back there willing to accept the package. Such a machine will be able to watch a motion picture, then decide how good it is and write an astute and insightful review of the movie.

But will this machine actually *enjoy* the movie? And, just as important, will *we* be able to know if it does? Here we come inevitably to the looming great challenge, and great puzzle, of this coming epoch: machine consciousness. Machines probably won't need consciousness to outperform us in almost every measurable way. Nevertheless, deep down we will surely regard them with a kind of disdain if they don't have it.

Trying to create consciousness may turn out to be the way we finally begin to understand this most deeply mysterious and precious of all human attributes. We don't understand how conscious experience arises or its purpose in human beings–why we delight in the sight of a sunset, why we are stirred by the *Eroica* symphony, why we fall in love. And yet, consciousness is the most remarkable thing the universe has ever created. If we, too, manage to create it, it would be humankind's supreme technological achievement, a kind of miracle that would fundamentally alter our relationship with our machines, our image of ourselves, and the future of our civilization. –GLENN ZORPETTE

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/spectrallines0617

ILLUSTRATION BY **Chad Hagen**

NEWS

€

€1 BILLION: BUDGET FOR PLANNED EU INITIATIVE TO ADVANCE QUANTUM TECHNOLOGIES

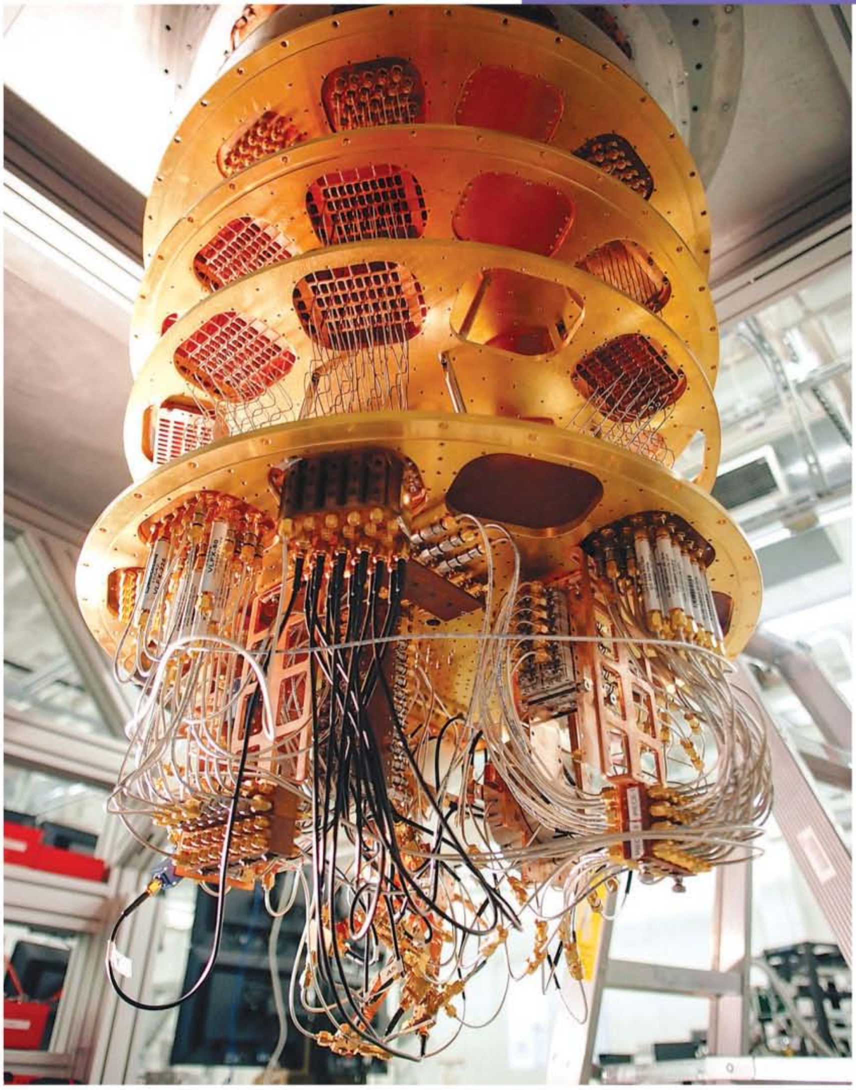

ERIK LUCERO

PUT CHIP HERE: Google will put its superconducting quantum computer chip in this 10-millikelvin dilution refrigerator.

GOOGLE AIMS FOR QUANTUM COMPUTING SUPREMACY

The company plans to build a 49-qubit chip to prove quantum computing's prowess

Quantum computers have long held the promise of performing certain calculations that are impossible—or at least, entirely impractical—for even the most powerful conventional computers to perform. Now, researchers at a Google laboratory in Goleta, Calif., may finally be on the cusp of proving it, using the same kinds of quantum bits, or qubits, that one day could make up large-scale quantum machines.

By the end of this year, the team aims to increase the number of superconducting qubits it builds on integrated circuits to create a 7-by-7 array. With this quantum IC, the Google researchers aim to perform operations at the edge of what's possible with even the best supercomputers, and so demonstrate "quantum supremacy."

"We've been talking about, for many years now, how a quantum processor could be powerful because of the way that quantum mechanics works, but we want to specifically demonstrate it," says team member John Martinis, a professor at

»

the University of California, Santa Barbara, who joined Google in 2014.

A system size of 49 superconducting qubits is still far away from what physicists think will be needed to perform the sorts of computations that have long motivated quantum computing research. One of those is Shor’s algorithm, a computational scheme that would enable a quantum computer to quickly factor very large numbers and thus crack one of the foundational components of modern cryptography. In a recent commentary in *Nature*, Martinis and colleagues estimated that a 100-million-qubit system would be needed to factor a 2,000-bit number–a not-uncommon public key length–in one day. Most of those qubits would be used to create the special quantum states that would be needed to perform the computation and to correct errors, creating a mere thousand or so stable “logical qubits” from thousands of less stable physical components, Martinis says.

There will be no such extra infrastructure in this 49-qubit system, which means a different computation must be performed to establish supremacy. To demonstrate the chip’s superiority over conventional computers, the Google team will execute operations on the array that will cause it to evolve chaotically and produce what looks like a random output. Classical machines can simulate this output for smaller systems. In April, for example, Lawrence Berkeley National Laboratory reported that its 29-petaflop supercomputer, Cori, had simulated the output of 45 qubits. But 49 qubits would push–if not exceed–the limits of conventional supercomputers.

This computation does not as yet have a clear practical application. But Martinis says there are reasons beyond demonstrating quantum supremacy to pursue this approach. The qubits used to make the 49-qubit array can also be used to make larger “universal” quantum systems with error correction, the sort that could do things like decryption, so the chip should provide useful validation data.

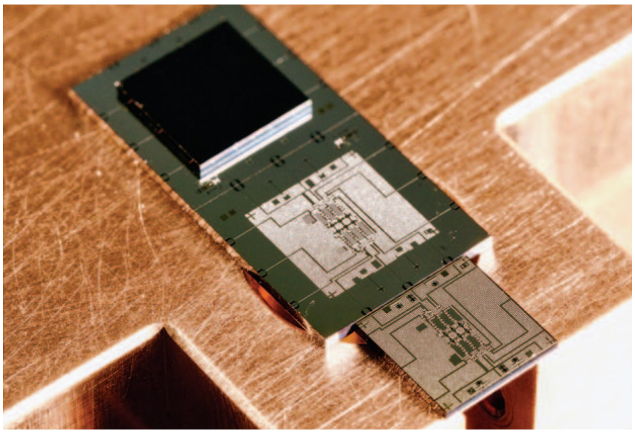

STEPS TO SUPREMACY: Google’s quantum computing chip is a 2-by-3 array of qubits. The company hopes to make a 7-by-7 array later this year.

ERIK LUCERO

There may also be, the team suspects, untapped computational potential in systems with little or no error correction. “It would be wonderful if this were true, because then we could have useful products right away instead of waiting for a long time,” says Martinis. One potential application, the team suggests, could be in the simulation of chemical reactions and materials.

Google recently performed a dry run of the approach on a 9-by-1 array of qubits and tested out some fabrication technology on a 2-by-3 array. Scaling up the number of qubits will happen in stages. “This is a challenging system engineering problem,” Martinis says. “We have to scale it up, but the qubits still have to work well. We can’t have any loss in fidelity, any increase in error rates, and I would say error rates and scaling tend to kind of compete against each other.” Still, he says, the team thinks there could be a way to scale up systems well past 50 qubits even without error correction.

Google is not the only company working on building larger quantum systems without error correction. In March, IBM unveiled a plan to create such a superconducting qubit system in the next few years, also with roughly 50 qubits, and to make it accessible on the cloud. “Fifty is a magic number,” says Bob Sutor, IBM’s vice president for this area, because that’s around the point where quantum computers will start to outstrip classical computers for certain tasks.

The quality of superconducting qubits has advanced a lot over the years since D-Wave Systems began offering commercial quantum computers, says Scott Aaronson, a professor of computer science at the University of Texas at Austin. D-Wave, based in Burnaby, B.C., Canada, has claimed that its systems offer a speedup over conventional machines, but Aaronson says there has been no convincing demonstration of that. Google, he says, is clearly aiming for a demonstration of quantum supremacy that is “not something you’ll have to squint and argue about.”

It’s still unclear whether there are useful tasks a 50-or-so-qubit chip could perform, Aaronson says. Nor is it certain whether systems can be made bigger without error correction. But he says quantum supremacy will be an important milestone nonetheless, one that is a natural offshoot of the effort to make large-scale, universal quantum machines: “I think that it is absolutely worth just establishing as clearly as we can that the world does work this way. Certainly, if we can do it as a spin-off of technology that will be useful eventually in its own right, then why the hell not?”

–RACHEL COURTLAND

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/supremacy0617

ISTOCKPHOTO

NEW BUZZ ON BRAIN STIMULATION FOR DEPRESSION

Treatments that target specific neural circuits provide better results

Here are three things we know for sure about deep brain stimulation (DBS) as a treatment for severe depression:

1. When the pacemaker-like brain implants do help depressed people, those people get dramatically better. “I have patients who got their implants 10 years ago now,” says Helen Mayberg, a professor of psychiatry and neurology at Emory University, in Atlanta. “These people get well and they stay well,” she says.

2. Unfortunately, DBS doesn’t help everybody. Experimental trials by Mayberg and others have consistently had a subset of people who simply don’t respond to the treatment. And two big industry-sponsored trials were counted as failures by the companies involved, squashing hopes that the treatment would soon be available for mainstream clinical use.

3. Researchers must figure out why some depressed patients respond to DBS while others don’t; otherwise their experiments will never lead to a truly practical treatment that can gain regulators’ approval. Two new academic studies show that researchers are progressing toward that understanding.

The treatment, which requires brain surgery for the implantation of electrodes and then constant pulses of stimulation, is based on a successful DBS treatment for Parkinson’s disease, which improves patients’ tremors and other movement problems. While the exact mechanism of DBS’s action in the brain is still unclear, many think the steady stimulating pulses override faulty patterns of electrical activity in the brain region that surrounds the electrodes.

For depression, much research and debate have centered on finding the proper target for stimulation. But Mayberg now advocates thinking of depression “in neural circuit terms.” She has long focused on a brain region called the subcallosal cingulate (more specifically, on a piece of the SCC named Brodmann’s area 25), which is known to play a role in regulating mood. Recently her team has also begun scrutinizing the bundles of neural fibers that project from the SCC to connect to other brain regions. “We have to rethink things: It’s not about the target but the projections,” Mayberg says.

In a study reported in the journal *Molecular Psychiatry* in April, Mayberg’s team treated 11 depression patients who hadn’t responded to drugs, talk therapy, or electroconvulsive therapy. Before the surgery, the researchers used MRI scans of each patient to identify the exact location of four fiber bundles that had been identified as crucial in 2014 research involving patients who responded to DBS. That earlier study looked for commonalities in responders’ brains, and it found they all had strong connections along those four tracts. “We can now use that template,” Mayberg says.

Surgeons working with Mayberg implanted the electrodes in a location that ensured that the pulses of electricity would travel along those neural fibers, thus activating not only the SCC but also the connecting brain regions. When the 11 patients were evaluated after six months, eight of them were judged to be responding to treatment.

To try to help the nonresponders, the researchers evaluated the placement of their electrodes and their stimulation parameters. For one patient, the researchers adjusted the parameters to more closely match the desired template of fiber activation. "We changed it to get those fiber tracts stimulated, and he got better," Mayberg says. For the remaining two nonresponders, there was no clear anatomical explanation for their lack of response.

Another study, published in the journal *Brain Stimulation*, also stresses the importance of understanding the brain's connectome, which can be thought of as a circuit diagram of the brain. Thomas Schlaepfer, a professor of psychiatry at the University of Freiburg, in Germany, treats depression through a different neural circuit than Mayberg's group. His team implanted electrodes in the medial forebrain bundle, which is part of the neural circuitry involved in feelings of pleasure and motivation. After one year, six of his eight patients had responded positively to the DBS treatment.

> "I was never discouraged by the failure of the industry-sponsored clinical trials.... I think there will be a place for DBS in the treatment of depression"
>
> —Herbert Ward, University of Florida

For one of the nonresponders, Schlaepfer says he saw a clear cause: The patient had experienced a hemorrhage during the implantation surgery. "When we did imaging, we saw that the hemorrhage was at exactly the site where we stimulated," he says. "That probably robbed her of the ability to respond, which was very tragic for her—but almost proved the point."

With two different sets of neural circuits under investigation and showing promising results, researchers can imagine a day when depressed patients go in for brain imaging so doctors can map their individual circuitries and devise custom-made treatment plans. "No treatment works for everyone," Mayberg says. "We should figure out who does best with what. The state of these brain networks will determine how best to intervene."

Herbert Ward, a psychiatrist at University of Florida Health and an expert on DBS, calls both Mayberg and Schlaepfer's results "very promising" and expects them to guide future work to prove DBS's efficacy. "I was never discouraged by the failure of the industry-sponsored clinical trials," Ward says, which targeted brain structures rather than their connections. If future trials focus on the circuitry of depression instead, Ward believes the treatment will eventually gain regulatory approval. "I think there will be a place for DBS in the treatment of depression," he says.

—ELIZA STRICKLAND

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/dbsdepression0617

THE MOST COMPLEX 2D MICROCHIP YET

This molybdenum disulfide microprocessor contains more than 100 transistors

Scientists hope that two-dimensional materials such as graphene or molybdenum disulfide will allow Moore's Law to continue apace once it becomes impossible to make further progress using silicon. A 3-atom-thick microchip developed by researchers at the Vienna University of Technology (VUT) may be the first example of 2D materials following the seemingly inexorable growth in the number of transistors in integrated circuits that Gordon Moore observed decades ago. Previously, the number of transistors on ICs made from 2D materials had remained in the single digits, says study senior author Thomas Müller, an electrical engineer at VUT. The chip that he and his colleagues created boasts 115 transistors. They described their device on 11 April in the journal *Nature*.

The microchip can execute user-defined programs stored in external memory, perform logic operations, and transmit data to its periphery. Although this prototype operates on single-bit data, the researchers say their design is readily scalable to multibit data. They also note their invention is compatible with existing semiconductor manufacturing processes.

The new device is made of a thin film of molybdenum disulfide, which resembles a sheet of molybdenum atoms sandwiched between two layers of sulfur atoms. The film is only six-tenths of a

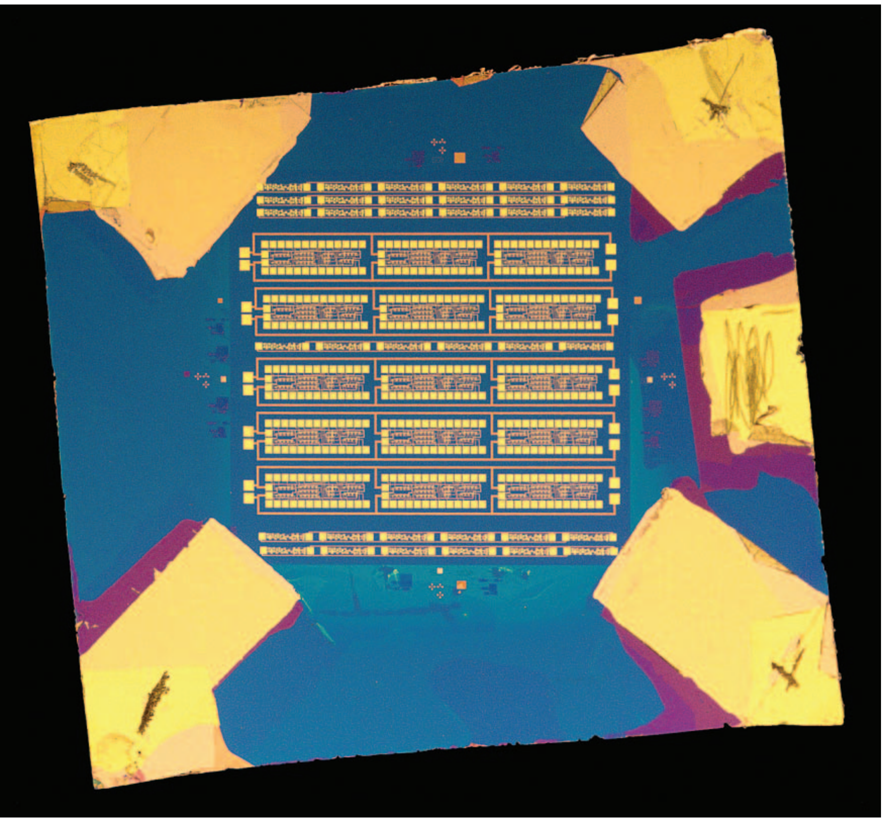

FLAT AND PACKED: The prototype processor consumes about 60 microwatts and contains 115 transistors. Its designers hope to integrate more.

STEFAN WACHTER

nanometer thick. In comparison, the active layer of a silicon microchip can be up to about 100 nanometers thick.

"It's exciting that the authors have been able to build circuits of this complexity using an atomically thin 2D material that has only been studied for five years," says Eric Pop, an electrical engineer at Stanford University, who did not take part in this research.

So, what's molybdenum disulfide's Moore's Law potential? The minimum size of features in the VUT prototype was a rather chunky 2 micrometers. However, "Going to 200- or 100-nanometer transistor channel lengths should be rather straightforward," says Müller. He adds that improvements in the quality of electrical contacts in these circuits should result in an ultimate scaling limit for 2D transistors of about 1 nanometer. "This cannot be reached with silicon; the limit there is around 5 nanometers," says Müller.

The main obstacle to building even more complex microchips with molybdenum disulfide is the trouble the team has had with fabricating transistors. Currently, Müller's team's yield for fully functional chips is only a few percent. One way to increase transistor yield is to grow more-uniform molybdenum disulfide films, he says. Unfortunately, imperfections in the films crop up when they are transferred from the sapphire substrates on which they are grown to their target wafers, he explains. The researchers are now working on ways to grow the molybdenum disulfide film directly on the target wafer to remove this transfer step. By improving the uniformity of molybdenum disulfide films, "it should be rather straightforward to increase the complexity of two-dimensional circuits to tens of thousands of transistors," Müller says.

The target wafer for the prototype device was silicon, but the researchers say it could, in principle, be produced on virtually any backing. "If circuits could be demonstrated on flexible substrates doing a unique function that silicon circuits cannot do, that would be an important future development," Stanford's Pop says.

But another important change is necessary if 2D microprocessors are going to incorporate hundreds of millions of transistors, as modern silicon chips do. Engineers will have to be able to switch from the *n*-type transistor design used in this prototype to the lower-power CMOS designs used in conventional microchips today, Müller says. "This probably will require another two-dimensional semiconductor material instead of molybdenum disulfide, but there are plenty—for example, tungsten diselenide," he says.

As for performance, the total power consumption of the prototype circuit is roughly 60 microwatts, and it operates at frequencies between 2 and 20 kilohertz, say the VUT researchers. "Our device is, of course, by no means competitive with current silicon-based microprocessors. It is just a very first step towards a new generation of electronic devices," Müller acknowledges.

Stanford's Pop agrees with that assessment. The VUT chip's field-effect mobility—the ease with which electrons can flow through transistors—when compared with other molybdenum disulfide films is "at least an order of magnitude worse than the present state of the art in the published literature from many groups, including ours," he explains. That mobility will have to improve, Pop says. "Molybdenum disulfide circuits with low mobility will operate slowly and inefficiently."

—CHARLES Q. CHOI

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/2dprocessor0617

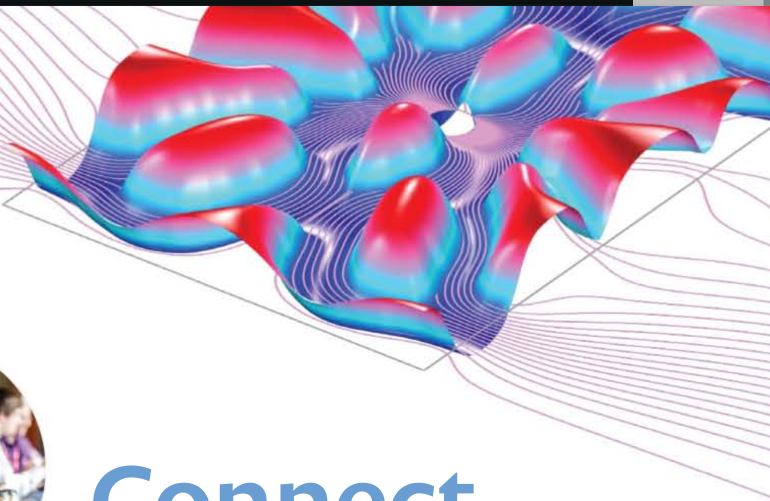

RESOURCES

1961: THE YEAR THE ZUSE GRAPHOMAT Z64 PLOTTER WAS INTRODUCED. IT WAS SOON ADOPTED BY SOME OF THE FIRST DIGITAL ARTISTS.

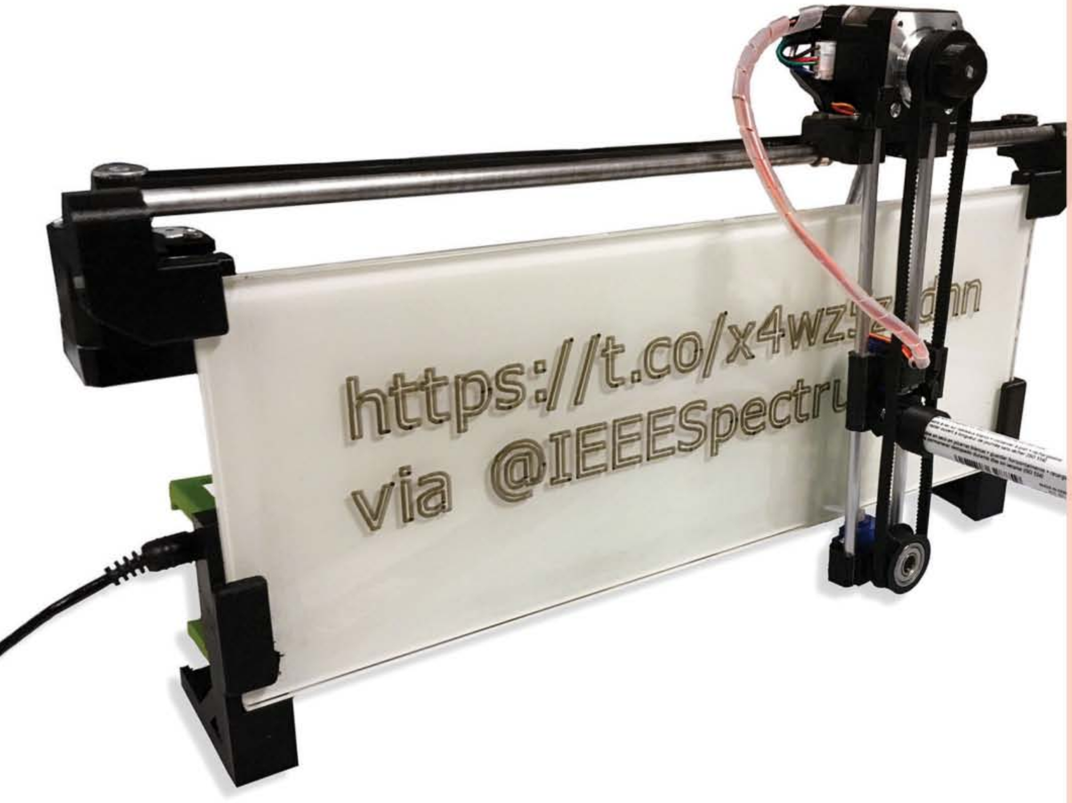

THE MASTER PLOTTER

THE iBOARDBOT MAKES A FUN DISPLAY FOR VECTOR-BASED IMAGES OR TEXT

RESOURCES_HANDS ON

X-Y plotters were once a common sight in places that needed high-quality hard copies of computer-generated line drawings. Blueprints and charts of all types could be produced with multiple colors on large pieces of paper, with perfectly smooth lines and curves. But eventually pixel-based laser and inkjet printers increased sufficiently in resolution—and decreased sufficiently in price—to make plotters something of a rarity. ● Still, some things were lost in this particular march of progress. One casualty was the distinctive lines of vector-based graphics and text. Another was *fun*: Watching a laser printer print is about as interesting as watching a refrigerator, while there's something mesmerizing in seeing the head of a plotter spring to life and construct an image line by line through confident sweeps and arcs. The US $171 iBoardbot from JJRobots is intended to bring back some of the romance of the plotter, but with a contemporary twist—it's cloud controlled and self-erasable. ● The iBoardbot comes either as a kit or in a largely assembled version for $234. JJRobots sent me the kit, which has four major sets of components. First is the erasable drawing surface itself, made up of a 40- by 15-centimeter piece of tempered glass (you can choose from multiple colors). Second is the electronics, consisting of an Arduino and a custom shield created by JJRobots to drive the motors. The motors—two stepper motors to control the X and Y positions of the plotter head, along with two servos to handle the raising and lowering of the plotter and wipe heads—are the third group. And finally, there's the collection of metal rods and plastic widgets that hold the whole thing together. ● On its website, JJRobots estimates that all these components can be assembled into a working iBoardbot ▸

RANDI KLETT

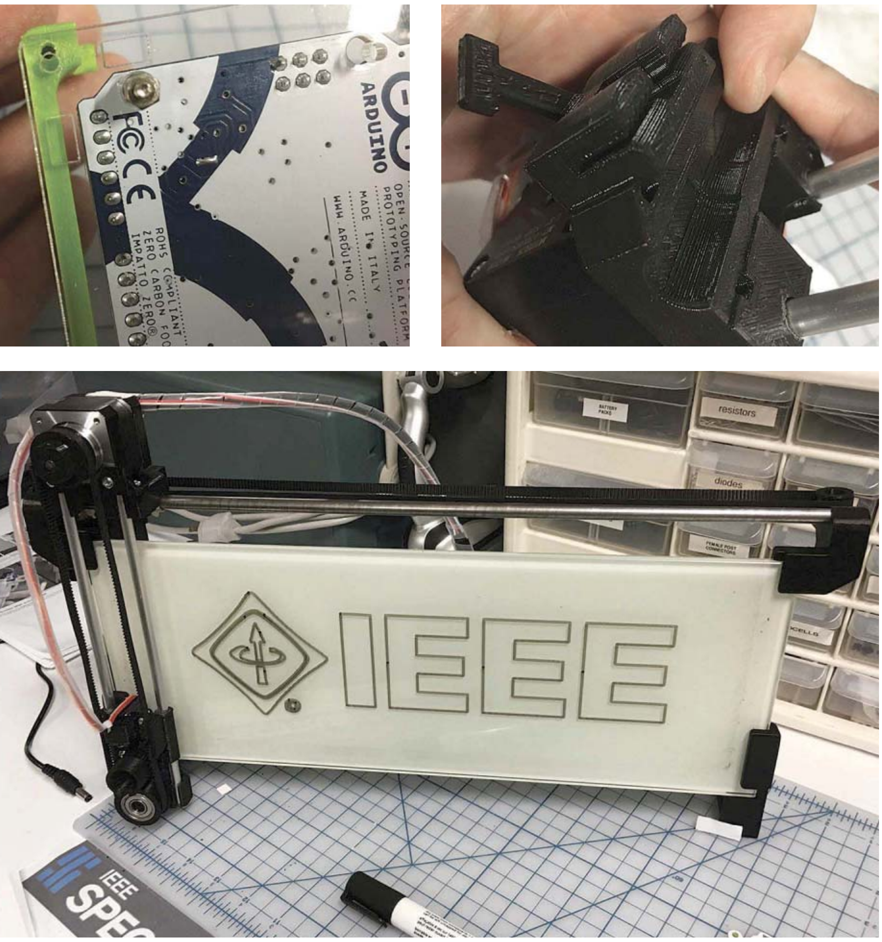

PRINTER TO PLOTTER: Some components of the iBoardBot are 3D printed [top right]. I had some alignment problems [top left] but got nice text and images nonetheless [bottom].

in “an hour or so,” which—no. I have no doubt that if I built a bunch of these things, I could get the time down to an hour. For someone unfamiliar with the kit, simply identifying which oddly shaped piece of plastic is the one called for in the next step of the instructions (and then making sure it’s the right way around) can take a few moments. But for me, the biggest time killer was that the plastic widgets are 3D printed. They’re complex shapes, so it makes sense to print them, but I did have problems with some of the widgets being just a tiny bit smaller than they should be. (Items that are just a tiny bit too large or too small are a recurring problem that can crop up in 3D printing.) This resulted in some mismatches between the printed parts and components made out of other materials.

In particular, one piece had two holes intended to grip a pair of metal rods. The plotter head uses these rods as tracks when moving up and down. After considerable time spent trying to twist the rods into the holes, as per the instructions, I realized the holes were simply too small. I used a cylindrical needle file to widen the holes just enough to accommodate the rods and still provide enough grip to hold the metal in place. The rest of the assembly went fairly smoothly, although I did come across another mismatch in size between the printed parts used to form the sides of a box used to hold the electronics and the acrylic plates that form the top and bottom of the box.

Once all was assembled, the next step was to get the iBoardbot ready to accept instructions. The custom shield that drives the motors also has a Wi-Fi chip, and wireless is the only way to communicate with the system out of the box (although all the source code is available, so you could program it to accept commands over a USB cable if you really wanted to). The first step is to visit the JJRobots website and run its connection wizard using a unique key provided for each iBoardbot. Midway through the process, you connect your laptop or smartphone to the device, which initially acts as a standalone Wi-Fi access point. You then feed it the SSID and password details for your regular Internet-connected local Wi-Fi network.

Once that’s done, the drawing surface is controlled via a Web app hosted by JJRobots. Access is obtained by passing your key in the app’s URL, which is bad from a security point of view: If someone can get your key (a 16-character alphanumeric sequence), he or she can write to your board. The app allows you to send short text messages or graphics to the board. The graphics modes will let you either make freehand sketches in real time (it’s great fun to watch the board follow along with you as you draw on a touch screen) or upload an image—you can use scalable vector graphics files, or the Web app can do its best to turn a bitmapped photo into a vector image.

My initial attempts at getting text and images out of the iBoardbot required some trial and error—the servos controlling the pen and the wiper have to be set just so. This involved partially disassembling the iBoardbot, unscrewing, adjusting, and rescrewing the arms attached to the servos, and repeating until the pen and wiper were moving correctly into position.

But the most useful thing you can do is control the iBoardbot using the IFTTT (“if this, then that”) online service, which provides a number of applets that respond to certain triggers. For example, the iBoardbot can be programmed to act as a clock, writing out the time and date every hour. You can also write your own IFTTT “recipes.” Currently, my iBoardbot is set to write out tweets sent to or from the *IEEE Spectrum* twitter account, a far more entertaining way of monitoring our social media than something like TweetDeck. **—STEPHEN CASS**

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/handson0617

STEPHEN CASS (3)

EARNING YOUR DRONE WINGS

HERE'S HOW TO BECOME AN FAA-APPROVED "REMOTE PILOT"

HENRY SCHNEIDER

THE PILOT'S PROGRESS: Our intrepid airman, *IEEE Spectrum* senior editor David Schneider, can now take to the skies as an FAA-certified remote pilot.

Since 29 August 2016, the U.S. Federal Aviation Administration (FAA) has made it possible for people to fly small drones for commercial purposes, so long as they adhere to some generally reasonable rules described by Part 107 of Title 14 of the Code of Federal Regulations, better known as the small UAS (unmanned aircraft systems) rule. The key point of the new regs is that you have to become a qualified "remote pilot"—and get an official certificate to prove it.

I don't want to fly drones commercially, but I have enjoyed flying radio-controlled models using first-person view, or FPV, in which you fly by watching live video from a camera mounted on the drone, sometimes via virtual reality goggles. The FAA ostensibly forbade FPV drone flight in 2014 when it published its interpretation of what constitutes a model aircraft. The FAA contended that pilots have to be able to see the aircraft directly with their own eyes from the ground. Many modelers ignore that prohibition, a practice that the Academy of Model Aeronautics endorses. But I'm a play-it-by-the-rules sort of guy, so I put the hobby into mothballs. With the new regulations, though, I decided to get certificated as a bona fide remote pilot and dust off my FPV gear. After all, how hard could it be?

Having just earned my drone wings, I can now answer that question in some detail.

First off, it costs money: US $150, which pays for a proctored multiple-choice exam at an FAA-approved testing center. I used freely available study materials, starting with 3D Robotics' helpful guide. But some people might find boning up for the exam less painful by taking an online course, which can be had for about $150 on up. You'll also have to pass a Transportation Security Administration background check and fill out some online paperwork.

What sort of things does the FAA want to confirm that you know? Not, as it turns out, a lot of the technical nitty-gritty of operating a drone. Instead, the test focuses on things of greater significance to "real" pilots, like an understanding of different classes of air space and how to decipher the telegraphic weather reports for aviators.

There is also a slew of details to memorize about the FAA's regulations. Question: What special requirements come into play if you operate a drone during civil twilight? Answer: You need to use anticollision lights. Question: What do you do if you have to deviate from the rules in the interest of safety? Answer: Inform the FAA if the FAA requests it. (Figuring out why the FAA would ever ask to be informed about something it has no knowledge of remains an exercise for the student.) Question: How long do you have after permanently changing your mailing address to inform the FAA? Answer: 30 days.

I suspect that folks at the FAA constructed their remote-pilot exam with two thoughts in mind. First, they want to ensure that remote pilots are serious enough about earning the privilege to put in at the very least a solid weekend, if not a week or more, of study. Second, and more important, the FAA probably wants to instill into the newly certified the sense of joining the ranks of bona fide aviators and so make them more willing to embrace the norms of that community. Navigating things around the sky is a tradition that goes back more than a century, with its own culture, complete with quirky terms, acronyms, codes, concepts, and—above all—a shared commitment to be part of a complex, cooperative enterprise designed to keep everybody safe.

So I'm inclined to excuse the FAA for forcing me to memorize what each letter stands for in IMSAFE, PAVE, CTAF, AWOS, and TAF. I managed to do that without too much grief and scored a respectable 93 percent on the exam. (It takes only 70 percent to pass.) And who knows, one day the information might prove useful should I ever find myself on the TV game show "Jeopardy!": "I'll take Aviation Terms for $200, Alex!" **—DAVID SCHNEIDER**

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/remotepilot0617

HOME 3D PRINTER SHOWDOWN

THE STRIPPED-DOWN MINIMAKER AND THE MOD-T ARE DESIGNED TO BE KID FRIENDLY

THE EARLY HYPE THAT EVERY HOME WOULD SHORTLY have a 3D printer fizzled fast, but there are still companies interested in catering to the home (and classroom) market. Two—XYZprinting and New Matter—were brave enough to allow me to borrow their "kid friendly" models for several months to check them out. For XYZ, that's the da Vinci miniMaker; for New Matter, it's the Mod-t. These gadgets (both under US $300) are aimed at tweens and up, though the primary colors of the miniMaker seem designed to appeal to far younger children. ("Up" includes non-tech-savvy teachers and other adults interested in 3D printing who aren't hard-core makers.) ● Both of these school/home printers use only nontoxic PLA filament as a printing material, a biodegradable polyester made out of renewable resources, typically cornstarch. Although advanced 3D-printer users would likely see this as a limitation for the home market, it's a good thing; it means you don't have to worry about your kids using materials that require serious ventilation for safety. ● Neither printer takes up much room: The miniMaker measures 39 by 33.5 by 36 centimeters, while the Mod-t measures 38.3 by 34 by 29.5 cm. Their build volumes are about the same size: approximately 15 by 15 by 15 cm and 15 by 10 by 12.5 cm, respectively. ● I was pretty excited to try these printers out. I thought they would be fun to share with two of my admittedly pretty grown-up kids (ages 18 and 21) over the winter holidays. The 18-year-old had some experience with 3D printing and maker spaces; the 21-year-old had a number of things she wanted to print, such as a stand for her iPad and tchotchkes for sorority giveaways. ● I turned most of the unboxing and setup over to my 18-year-old as I watched. He started with the miniMaker (about $250). Setup took far longer than it should have—the instructions were cryptic at best and seemed to omit critical steps. We had some problems with the calibration test and had to repeat it a couple of times. ● By contrast, the unboxing and setup process for the Mod-t ($299) went smoothly, with clearly written instructions accompanied by impossible-to-misunderstand videos.

XYZPRINTING

However, disaster struck both machines: A firmware problem with the miniMaker and a busted motor on the Mod-t (my particular printer—a press "loaner"—had been through some rough handling previously) meant that my plans for a 3D-printer-themed holiday were over.

In late January, I was ready to start again—on my own this time. Given that I usually start dreaming about the beach around then, I decided to print various beach gadgets—like cup holders, towel pins, and sarong rings—using designs from the Thingiverse online repository.

The user interface and tools provided for the two systems differ in approach—mainly because the Mod-t requires you to upload designs to a website for processing, while the miniMaker does it on your computer—but both were straightforward.

I started with a large clothespin on the miniMaker; I planned to use the pin to clip

THE DUELING PRINTERS: The da Vinci miniMaker [left] and the Mod-t [right] are intended for younger and inexperienced users.

NEW MATTER

my beach towel to my chair. The printer slipped out of alignment, and I had to stop and restart. Interestingly, this was the very last alignment problem I had with that printer; every other project went smoothly. The Mod-t, assigned to print the cup holder, also had an alignment problem, and I had to stop the project and restart. This was also the last alignment problem I had with that printer and remains a mystery.

From then on, the miniMaker worked smoothly, although I wasn't thrilled that in order to switch colors I had to get out a screwdriver and fumble with changing a chip that goes inside the spool hub. You can use only filament purchased from XYZ for this printer; the company says that this ensures quality control, and to be fair, the price isn't outrageous, at around $28 for 600 grams of 1.75 mm filament. But the whole spool-and-chip disassembly process was annoying.

Meanwhile, the Mod-t was driving me a little bit crazy. The main problem is that the filament breaks. It was easy enough to reload—at least most of the time—but I lost count of how many times I tried to print, or was well into a project, when I noticed nothing was coming out of the print head and had to stop. The printer doesn't detect these breaks itself, so if you walk away, you can hear the thing chunking along for hours and think all is well, but nothing is actually printing. I also had a troubling experience with the print head getting stuck and shutting the whole thing down; tech support eventually walked me through the process of unsticking it manually. I was told that the head is supposed to react to resistance when it rises too high, but it doesn't always do that. (The company is working on a firmware update to address the issue.) After the stuck head problem, I was able to print a few more objects. Then a filament break that led to a filament jam finally pushed me to the end of my patience; I ran the unload sequence several times and then gave up, turning to the miniMaker for the last few objects I'd wanted to print.

The verdict: The Mod-t looks cool, and I loved to watch it calibrate itself—that process is really neat. And the filament is easy to change; that's a plus for people who don't want every object they print to be the same color. Meanwhile, the miniMaker was ridiculously hard to set up, it looks silly on my desktop, and the filament is more annoying to change…but, well, it's still working. And the filament didn't break—not even once.

—TEKLA S. PERRY

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/3dprinters0317

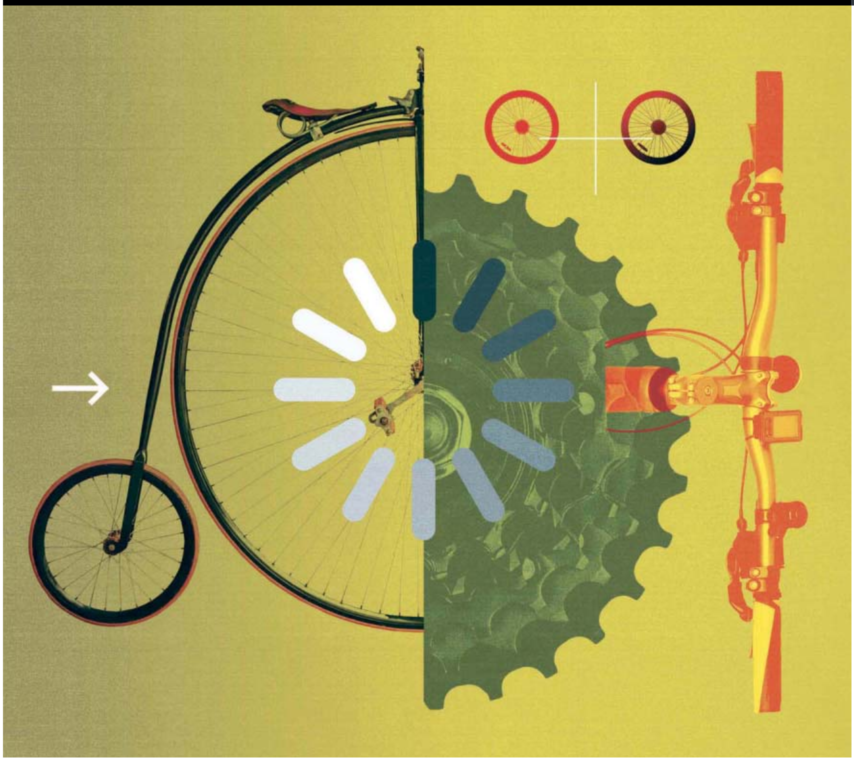

SLOW CYCLING

SOME TECHNICAL ADVANCES are delayed by either a failure of imagination or a concatenation of obstructive circumstances. I can think of no better example of both of these than the bicycle. ● Two centuries ago, on 12 June 1817 in Mannheim, Karl Drais, a forester in Germany's grand duchy of Baden, demonstrated for the first time his *Laufmaschine* ("running machine"), later also known as a draisine or hobby-horse. With the seat in the middle, front-wheel steering, and wheels of the same diameter, it was the archetype of all later vehicles that required constant balancing. However, it was propelled not by pedaling but by pushing one's feet against the ground, Fred Flintstone fashion. ● Drais covered nearly 16 kilometers (10 miles) in little more than an hour on his heavy wooden bicycle, faster than the typical horse-drawn carriage. But it's obvious, today at least, that the design was too heavy and clumsy and that there weren't yet enough suitable hard-top roads. But why, in the decades after 1820 that abounded with such inventions as locomotives, steamships, and manufacturing techniques, did it take so long to come up with a means of propulsion that could make the bicycle a practical machine, able to be ridden by anybody but infants? ● Only in 1866 did Pierre Lallement get his U.S. patent for a bicycle propelled by pedals attached to a slightly larger front wheel. Starting in 1868, Pierre Michaux made this *vélocipède* design popular in France. But the Michaudine did not become the precursor of modern bicycles; it was just an ephemeral novelty. The entire 1870s and the early 1880s were dominated by high-wheelers (also known as "ordinary" or penny-farthing bicycles) with pedals attached directly to the axles of front wheels with diameters of up to 1.5 meters (5 feet), to provide a longer distance per pedal revolution. These clumsy machines could be fast, but they were also difficult to mount and tricky to steer; their use called for dexterity, stamina, and a tolerance for dangerous falls.

Only in 1885 did two British inventors, John Kemp Starley and William Sutton, begin to offer their Rover safety bicycles with equally sized wheels, direct steering, a chain-and-sprocket drive, and a tubular-steel frame. Although it was not quite yet in the classic diamond shape, it was a truly modern bicycle design, ready for mass adoption. The trend accelerated in 1888, with the introduction of John Dunlop's pneumatic tires.

So a simple balancing machine consisting of two equally sized wheels, a minimal metal frame, and a short drive chain emerged more than a century *after* Watt's improved steam engines (1765), more than half a century *after* the introduction of mechanically far more complex locomotives (1829), years *after* the first commercial generation of electricity (1882)—but *concurrently* with the first designs of automobiles. The first light internal combustion engines were mounted in 1886 on three- or four-wheel carriages by Karl Benz, Gottlieb Daimler, and Wilhelm Maybach.

And although cars changed enormously between 1886 and 1976, bicycle design remained remarkably conservative. The first purpose-built mountain bikes came only in 1977. More widespread adoption of such novelties as expensive alloys, composite materials, strange-looking frames, solid wheels, and upturned handlebars began only during the 1980s.

Some inventions are tardy; others are precocious. One of my favorite examples of the latter is John Barber's 1791 British patent for "Obtaining and Applying Motive Power, & c. A Method of Rising Inflammable Air for the Purposes of Procuring Motion, and Facilitating Metallurgical Operations," which correctly outlined the operation of gas turbines. But in those days there were neither suitable steels for such a machine nor ways to generate the requisite power and pressure. The first working gas turbines came only in the late 1930s. ■

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/slowcycling0617

ILLUSTRATION BY **Stuart Bradford**

TECHNICALLY SPEAKING_BY PAUL McFEDRIES

OPINION

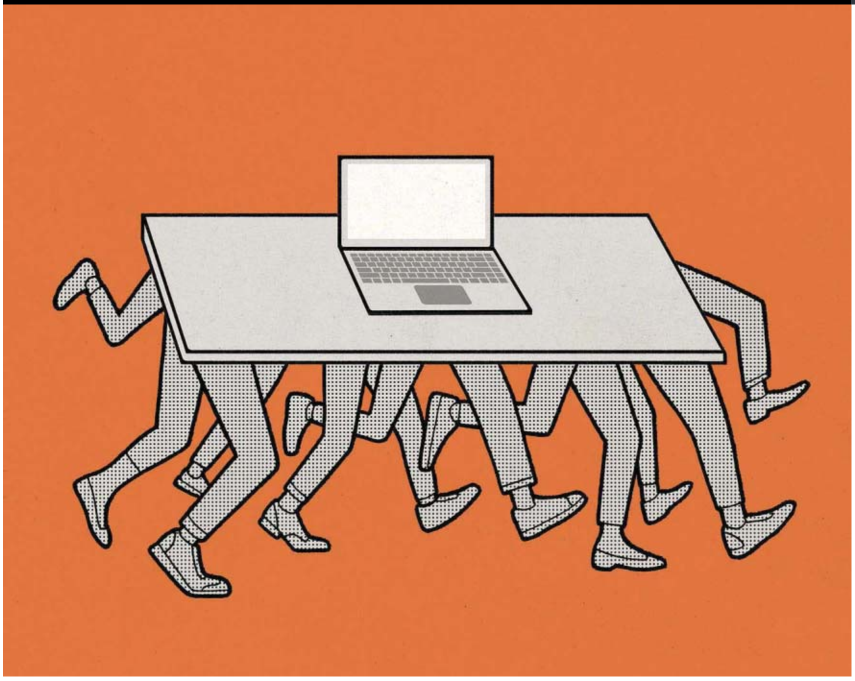

AGILE WORDS

> Through this work we have come to value: individuals and interactions over processes and tools; working software over comprehensive documentation; customer collaboration over contract negotiation; responding to change over following a plan. *—The Agile Manifesto*

IN HIS 1957 BOOK *Parkinson's Law, and Other Studies in Administration*, the naval historian and author C. Northcote Parkinson writes of a fictional committee meeting during which, after a two-and-a-half-minute nondiscussion on whether to build a nuclear reactor worth US $10 million, the members spend 45 minutes discussing the power plant's bike shed, worth $2,350. From this he coined Parkinson's Law of Triviality: "Time spent on any item of the agenda will be in inverse proportion to the sum involved." Using Parkinson's example, the programmer Poul-Henning Kamp popularized the term **bikeshedding**: frequent, detailed discussions on a minor issue conducted while major issues are being ignored or postponed. The functional opposite of bikeshedding is **trystorming**, which refers to rapidly and repeatedly prototyping or implementing new products and processes. In a bikeshedding culture, ideas get only a short discussion before being put off "for further study." In a trystorming culture, that same idea would be immediately prototyped, modeled, simulated, mocked up, or implemented, and examined to see what works and what doesn't. The trystorming motto is "Fail early, fail fast, fail often." ● That **fail-fast** mantra also describes the process of Agile software development, a set of principles for making software that emphasizes an iterative, collaborative, and adaptive approach. In one Agile framework called **Scrum**, programmers are assigned **timeboxes**–also known as **iterations** or **sprints**–which are set time periods in which they focus on completing a predefined goal, which then becomes a **WIP**: a work in progress. One of the key elements of Agile is that these goals are almost always usable features of the final product (a process called **continuous integration**). Programmers work in **Scrum teams** under the supervision of a **Scrum master**, who runs daily **standups**, short meetings to review progress in which people physically stand rather than sit.

Another Agile framework is **Kanban**, where remaining tasks are visualized for all to see (*kanban* is Japanese for "billboard"), and coders "pull" tasks from this list rather than having tasks "pushed" to them. What if Kanban teams want to use some aspects of the Scrum methodology, such as standups? Then call this hybrid approach **Scrumban.**

Kanban (and indeed all Agile methodologies) are very visual, with team leaders providing lots of big, visible charts depicting project metrics such as the **backlog** (showing the features remaining in the project) and either the **burndown** (the rate at which the backlog is being cleared) or the **burnup** (the rate at which project tasks are being completed). Some teams even use a **niko-niko calendar** (*nikoniko* is Japanese for "smile"), where at the end of each day team members display a visual indication of their mood (hence the alternative name **mood board**). All such charts and boards are often called **information radiators**.

Agile teams are big into **pair programming**, where two programmers sit at the same workstation, with one operating the keyboard and mouse and the other providing general direction and on-the-fly code reviews. These are known as the **driver** and the **navigator**, respectively, and they typically swap roles frequently. Both programmers are expected to provide running commentary throughout the process, which is known as **programming out loud.** A variation on that theme is **mob programming**, where the role of the navigator is assumed by the entire software team. A similar idea is **swarming**, where multiple team members jump into an in-progress task to help complete the work.

The main aim of Agile is to create happy customers by building great products. But look closer and you see a humane streak running through all of Agile. Teams seek a **sustainable pace,** meaning a work rate–sometimes called a **velocity**–that can be kept up indefinitely. ■

↗ **POST YOUR COMMENTS** at http://spectrum.ieee.org/technicallyspeaking0617

ILLUSTRATION BY Dan Page

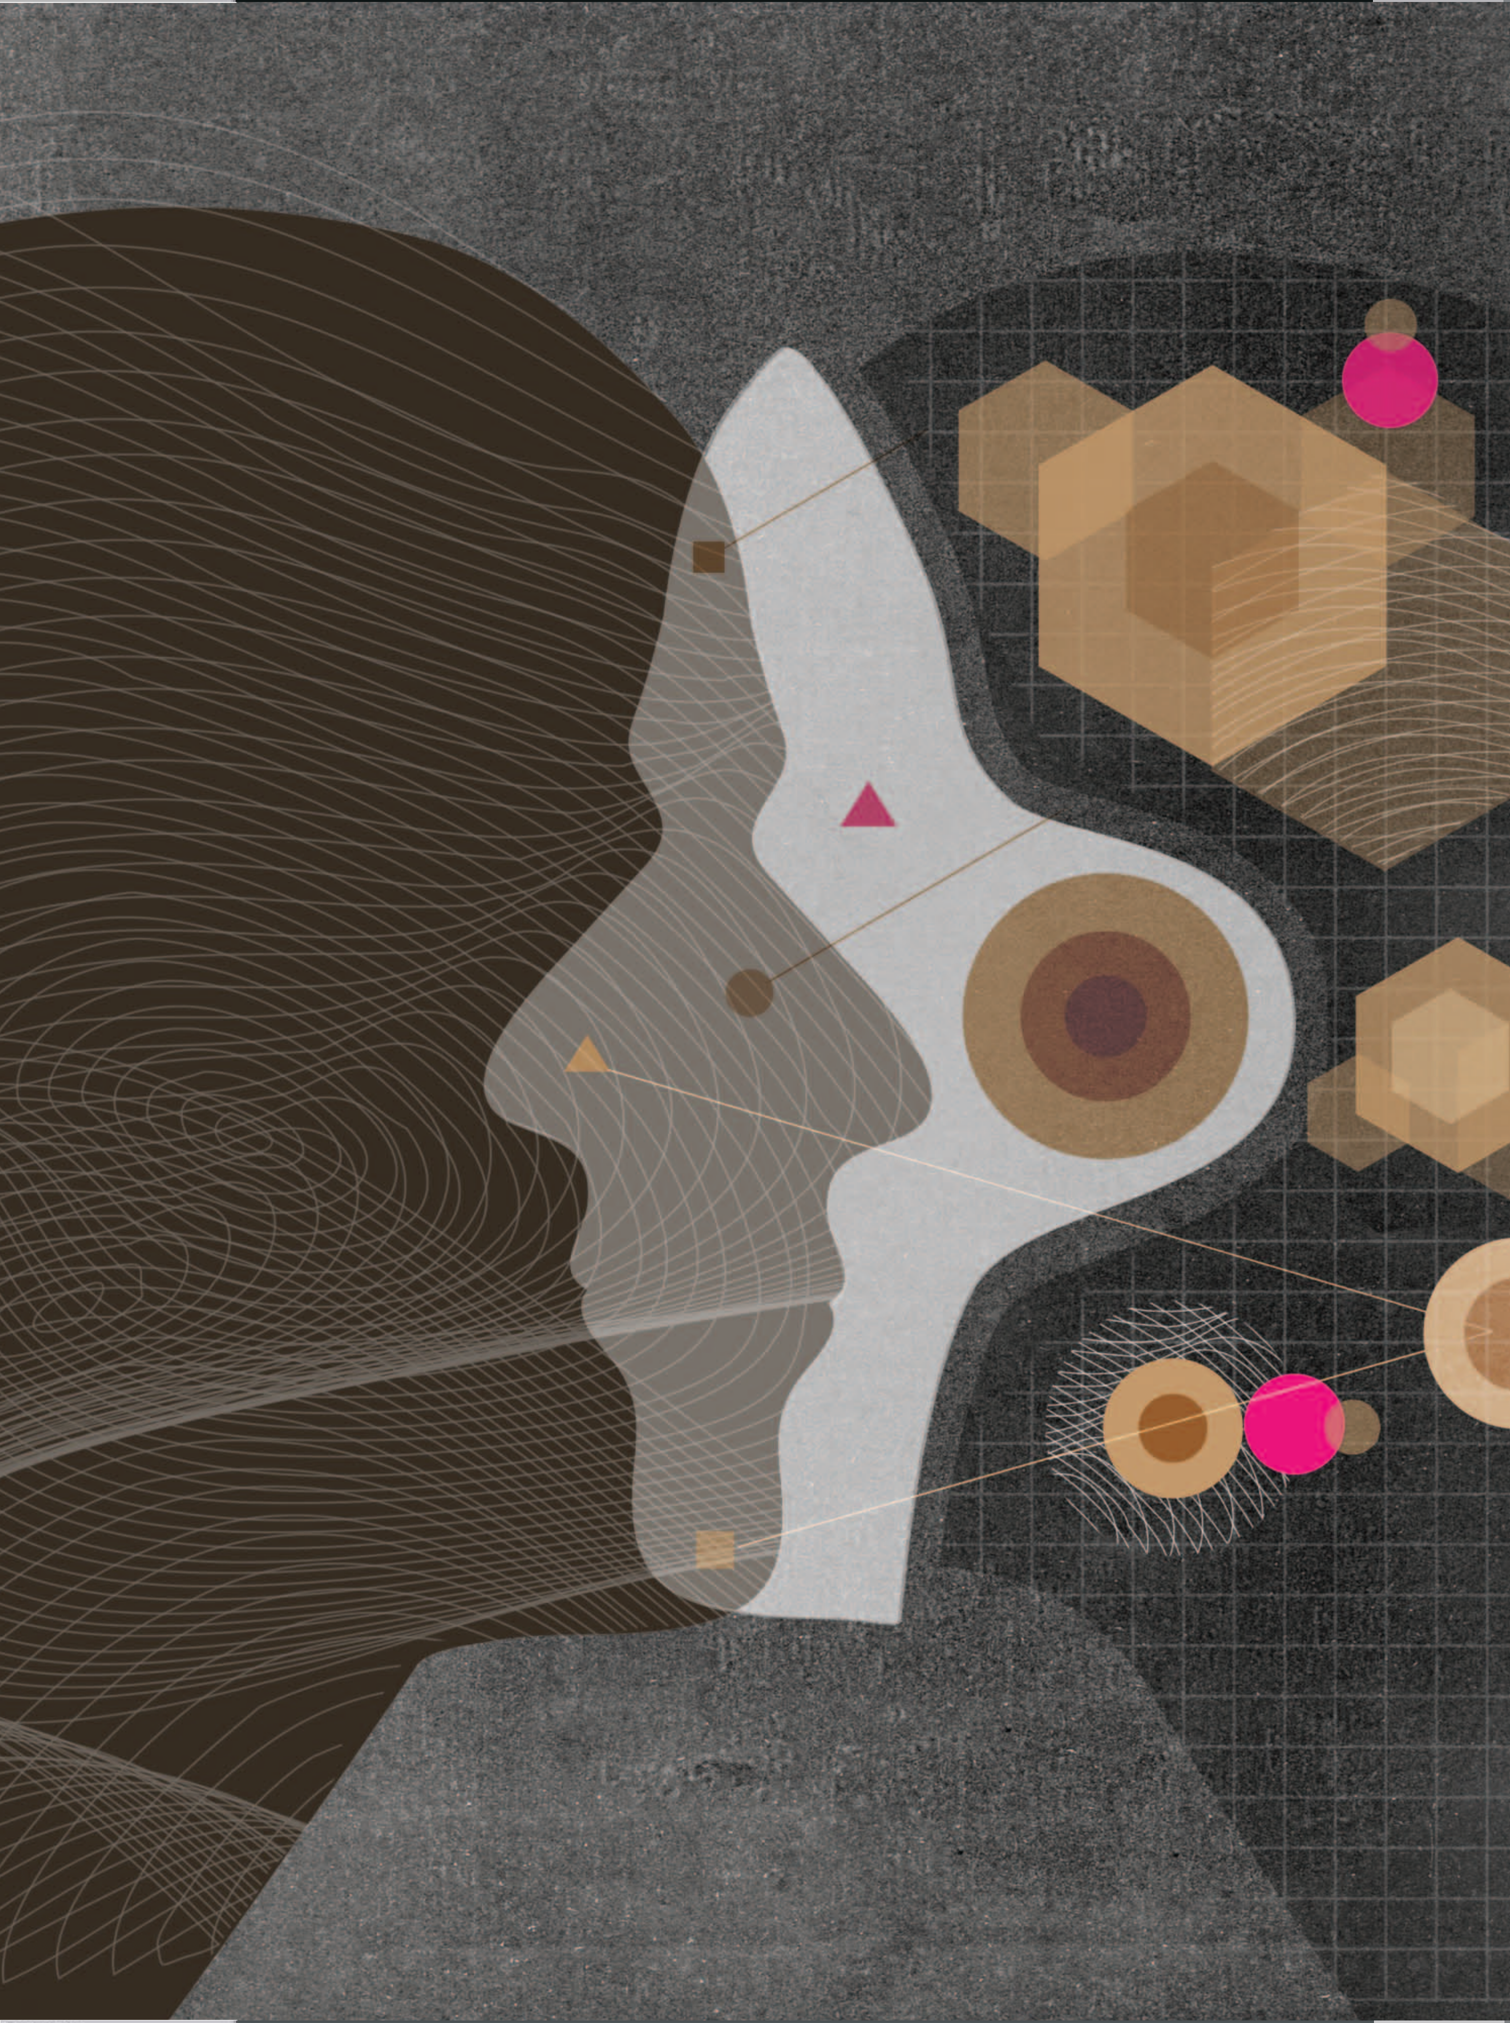

The Dawn of the Real Thinking Machine

NEUROSCIENCE WILL GIVE US WHAT WE'VE SOUGHT FOR DECADES: COMPUTERS THAT THINK LIKE WE DO

By **Fred Rothganger**

STEVE SITS UP and takes in the crisp new daylight pouring through the bedroom window. He looks down at his companion, still pretending to sleep. "Okay, Kiri, I'm up." ● She stirs out of bed and begins dressing. "You received 164 messages overnight. I answered all but one." ● In the bathroom, Steve stares at his disheveled self. "Fine, give it to me." ● "Your mother wants to know why you won't get a real girlfriend." ● He bursts out laughing. "Anything else?" ● "Your cholesterol is creeping up again. And there have been 15,712 attempts to hack my mind in the last hour." ● "Good grief! Can you identify the source?" ● "It's distributed. Mostly inducements to purchase a new RF oven. I'm shifting ciphers and restricting network traffic." ● "Okay. Let me know if you start hearing voices." Steve pauses. "Any good deals?" ● "One with remote control is in our price range. It has mostly good reviews."

ILLUSTRATION BY Chad Hagen

"You can buy it."

Kiri smiles. "I'll stay in bed and cook dinner with a thought."

Steve goes to the car and takes his seat.

Car, a creature of habit, pulls out and heads to work without any prodding.

Leaning his head back, Steve watches the world go by. Screw the news. He'll read it later.

Car deposits Steve in front of his office building and then searches for a parking spot.

Steve walks to the lounge, grabs a roll and some coffee. His coworkers drift in and chat for hours. They try to find some inspiration for a new movie script. AI-generated art is flawless in execution, even in depth of story, but somehow it doesn't resonate well with humans, much as one generation's music does not always appeal to the next. AIs simply don't share the human condition.

But maybe they could if they experienced the world through a body. That's the whole point of the experiment with Kiri....

THE FUTURE according to...

Robin Hanson

Author of The Age of Em: Work, Love, and Life When Robots Rule the Earth

When will we have computers as capable as the brain?

Within a century or so if we manage to achieve brain emulations, two to four centuries if not.

How will brainlike computers change the world?

Computers as capable as brains will be cheaper than brains, and so displace brains in almost all job tasks. Since computers can be made fast in factories, the economy could grow much faster, perhaps doubling every month. Humans would have to retire, but they'd own most capital, which would double as fast as the economy. So, collectively, humans get rich very fast.

Do you have any qualms about a future in which computers have human-level (or greater) intelligence?

This is an enormous change, as big as that from foraging to farming, or from farming to industry. If you don't have qualms about a change that big, you really aren't paying attention.

• • •

IT'S SCI-FI NOW, but by midcentury we could be living in Steve and Kiri's world. Computing, after about 70 years, is at a momentous juncture. The old approaches, based on CMOS technology and the von Neumann architecture, are reaching their fundamental limits. Meanwhile, massive efforts around the world to understand the workings of the human brain are yielding new insights into one of the greatest scientific mysteries: the biological basis of human cognition.

The dream of a thinking machine—one like Kiri that reacts, plans, and reasons like a human—is as old as the computer age. In 1950, Alan Turing proposed to test whether machines can think, by comparing their conversation with that of humans. He predicted computers would pass his test by the year 2000. Computing pioneers such as John von Neumann also set out to imitate the brain. They had only the simplest notion of neurons, based on the work of neuroscientist Santiago Ramón y Cajal and others in the late 1800s. And the dream proved elusive, full of false starts and blind alleys. Even now, we have little idea how the tangible brain gives rise to the intangible experience of conscious thought.

Today, building a better model of the brain is the goal of major government efforts such as the BRAIN Initiative in the United States and the Human Brain Project in Europe, joined by private efforts such as those of the Allen Institute for Brain Science, in Seattle. Collectively, these initiatives involve hundreds of researchers and billions of dollars.

With systematic data collection and rigorous insights into the brain, a new generation of computer pioneers hopes to create truly thinking machines.

If they succeed, they will transform the human condition, just as the Industrial Revolution did 200 years ago. For nearly all of human history, we had to grow our own food and make things by hand. The Industrial Revolution unleashed vast stores of energy, allowing us to build, farm, travel, and communicate on a whole new scale. The AI revolution will take us one enormous leap further, freeing us from the need to control every detail of operating the machines that underlie modern civilization. And as a consequence of copying the brain, we will come to understand ourselves in a deeper, truer light. Perhaps the first benefits will be in mental health, organizational behavior, or even international relations.

Such machines will also improve our health in general. Imagine a device, whether a robot or your cellphone, that keeps your medical records. Combining this personalized data with a sophis-

PHOTO-ILLUSTRATION BY Gluekit

ticated model of all the pathways that regulate the human body, it could simulate scenarios and recommend healthy behaviors or medical actions tailored to you. A human doctor can correlate only a few variables at once, but such an app could consider thousands. It would be more effective and more personal than any physician.

Re-creating the processes of the brain will let us automate anything humans now do. Think about fast food. Just combine a neural controller chip that imitates the reasoning, intuitive, and mechanical-control powers of the brain with a few thousand dollars' worth of parts, and you have a short-order bot. You'd order a burger with your phone, and then drive up to retrieve your food from a building with no humans in it. Many other commercial facilities would be similarly human free.

That may sound horrifying, given how rigid computers are today. Ever call a customer service or technical support line, only to be forced through a frustrating series of automated menus by a pleasant canned voice asking you repeatedly to "press or say 3," at the end of which you've gotten nowhere? The charade creates human expectations, yet the machines frequently fail to deliver and can't even get angry when you scream at them. Thinking machines will sense your emotions, understand your goals, and actively help you achieve them. Rather than mechanically running through a fixed set of instructions, they will adjust as circumstances change.

That's because they'll be modeled on our brains, which are exquisitely adapted to navigating complex environments and working with other humans. With little conscious effort, we understand language and grasp shades of meaning and mood from the subtle cues of body language, facial expression, and tone of voice. And the brain does all that while consuming astonishingly little energy.

That 1.3-kilogram lump of neural tissue you carry around in your head accounts for about 20 percent of your body's metabolism. Thus, with an average basal metabolism of 100 watts, each of us is equipped with the biological equivalent of a 20-W supercomputer. Even today's most powerful computers, running at 20 million W, can't come close to matching the brain.

How does the brain do it? It's not that neurons are so much more efficient than transistors. In fact, when it comes to moving signals around, neurons have one-tenth the efficiency. It must be the organization of those neurons and their patterns of interaction, or "algorithms." The brain has relatively shallow but massively parallel networks. At every level, from deep inside cells to large brain regions, there are feedback loops that keep the system in balance and change it in response to activity from neighboring units. The ultimate feedback loop is through the muscles to the outside world and back through the senses.

Traditionally, neurons were viewed as units that collect thousands of inputs, transform them computationally, and then send signals downstream to other neurons via connections called synapses. But it turns out that this model is too simplistic; surprising computational power exists in every part of the system. Even a single synapse contains hundreds of different protein types having complex interactions. It's a molecular computer in its own right.

And there are hundreds of different types of neurons, each performing a special role in the neural circuitry. Most neurons communicate through physical contact, so they grow long skinny branches to find the right partner. Signals move along these branches via a chain of amplifiers. Ion pumps keep the neuron's cell membrane charged, like a battery. Signals travel as short sharp changes of voltage, called spikes, which ripple down the membrane.

The power of the brain goes beyond its internal connections, and includes its ability to communicate with other brains. Some animals form swarms or social groups, but only humans form deep hierarchies. This penchant, more than any unique cognitive ability, enables us to dominate the planet and construct objects of exquisite complexity. Collectively, though, we humans are capable of achieving truly great things.

Now we are combining machine intelligence along with our own. As our systems–industrial, technological, medical–grow in sophistication and complexity, so too must the intelligence that operates them. Eventually, our tools will think for themselves, perhaps even become conscious. Some people find this a scary prospect. If our tools think for themselves, they could turn against us. What if, instead, we create machines that love us?

● ■ ●

STEVE ARRIVES HOME full of dread. Inside, the place is pristinely clean. A delicious aroma wafts from the new oven. Kiri is on the back porch, working at an easel. He walks up behind her. "How was your day?"

"I made a new painting." She steps away to show him. The canvas contains a photo-perfect rendition of the yard, in oils.

"Um, it's nice."

"You're lying. I can tell from your biosensors."

"Listen, Kiri, I have to take you back to the lab. They say you've progressed as far as you can with me."

"I like it here. Please let me stay. I'll be anything you want."

"That's the problem. You try so hard to please me that you haven't found yourself."

Water trickles down her cheek. She wipes it and studies her moist hand. "You think all this is fake."

Steve takes Kiri in a tight embrace and holds her for a long time. He whispers, "I don't know." ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/thinkingmachine0617

Special Report:
CAN WE COPY THE BRAIN?

An Engineer's Guide to the Brain

IN THE HUMAN BRAIN, higher-level information processing occurs in the neocortex, neural tissue that forms the outer layer of the cerebral cortex. In its intricate folds, brain cells work together to interpret sensory information and to form thoughts and plans. The neocortex is divided into regions that take the lead on different types of processing. However, much of today's neuroscience research focuses on mapping the connectome: the neural connections between regions. —ELIZA STRICKLAND

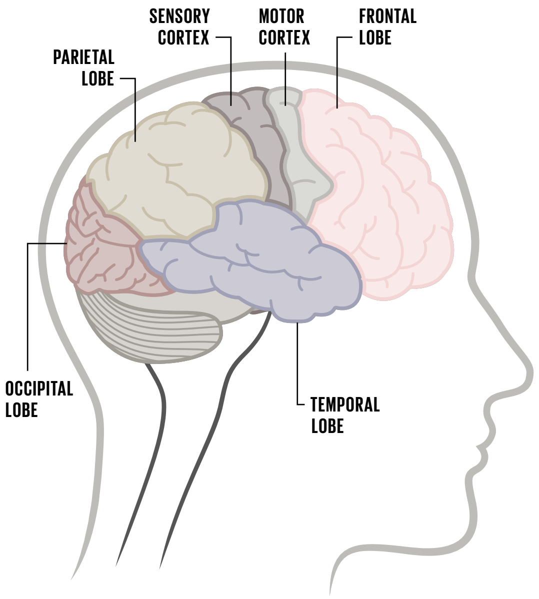

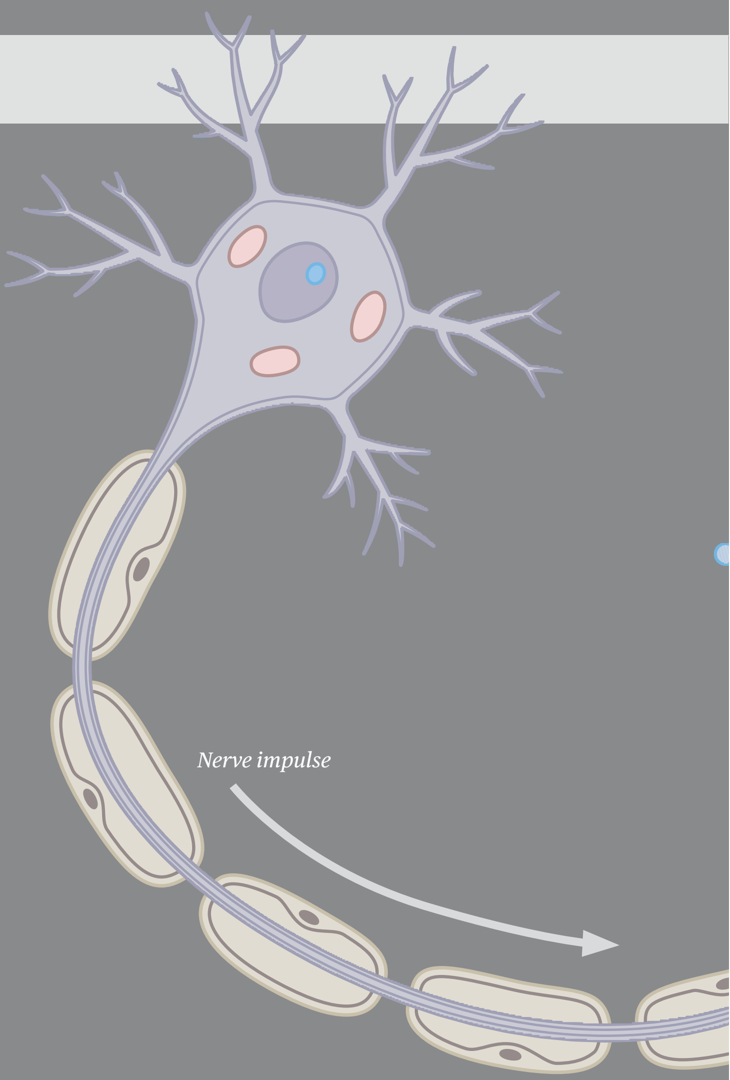

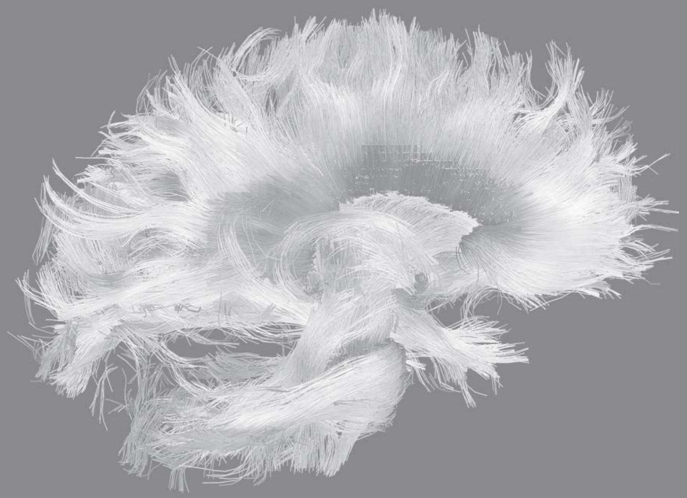

USC LABORATORY OF NEURO IMAGING/ATHINOULA A. MARTINOS CENTER FOR BIOMEDICAL IMAGING/CONSORTIUM OF THE HUMAN CONNECTOME PROJECT

1 A WIRING DIAGRAM

When neuroscientists trace the connections between brain regions, they focus on networks of electrically active brain cells called neurons. These neurons link up with one another, sending electrical signals through complicated circuits that span the brain. Each of the human brain's 86 billion neurons can connect to thousands of others.

2 NEURON TO NEURON

Each neuron receives incoming signals through branching structures called dendrites and outputs signals through its long axon. If the combination of incoming signals raises the neuron's voltage past a certain threshold, voltage-gated ion channels all along the axon begin to open.

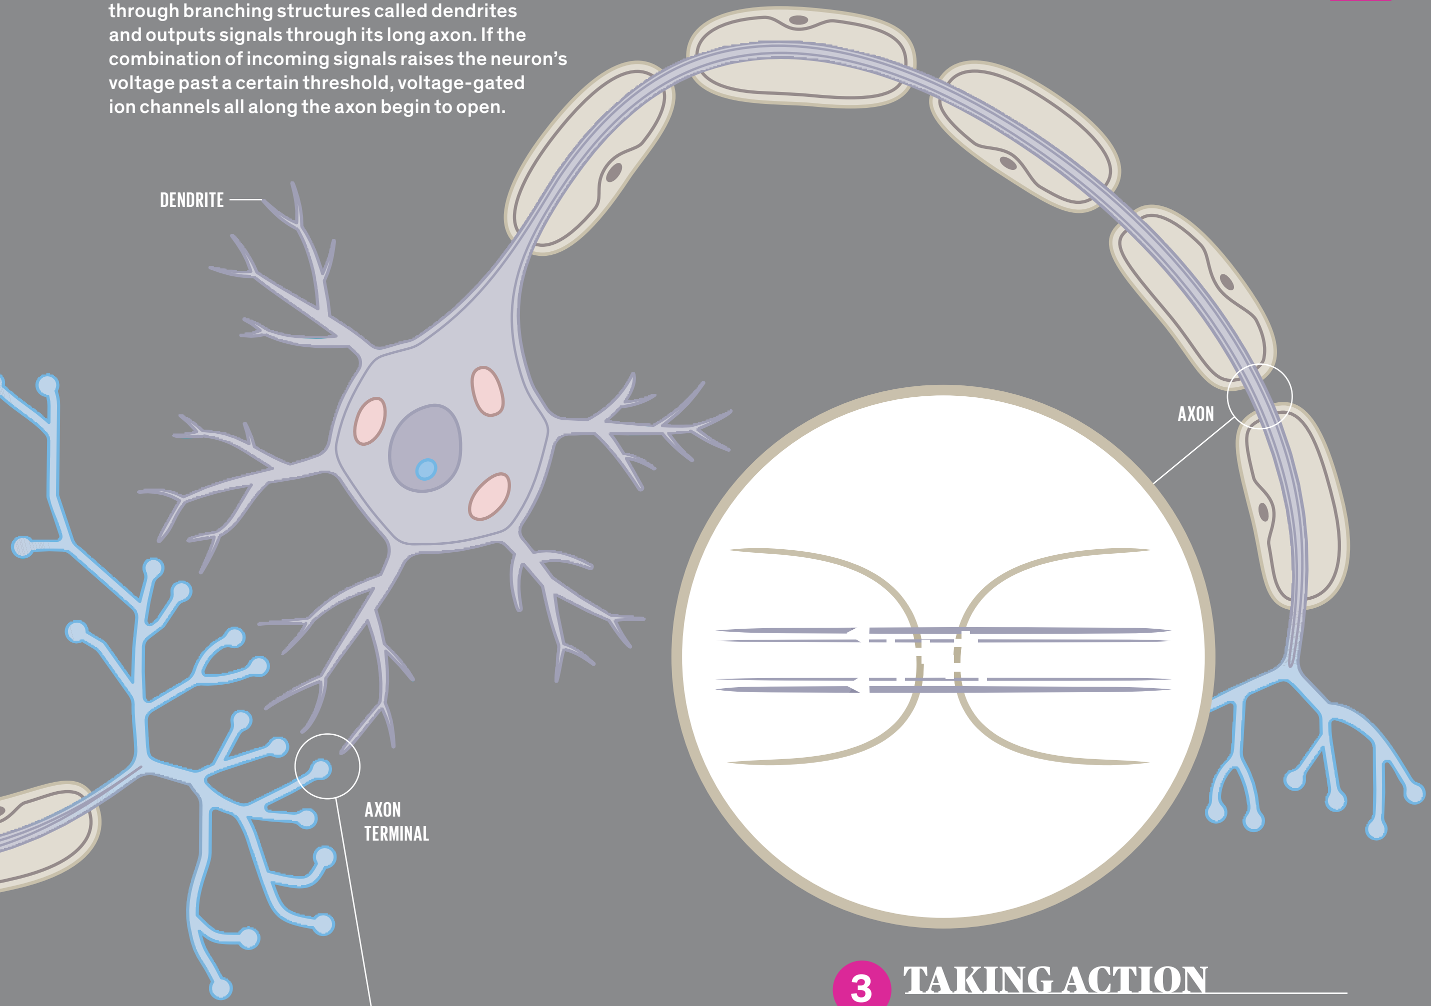

3 TAKING ACTION

The open channels allow ions to flow in, creating a propagating signal called an action potential that flows down the long axon to a gap called the synapse. There, the action potential triggers a signal to the connecting neurons.

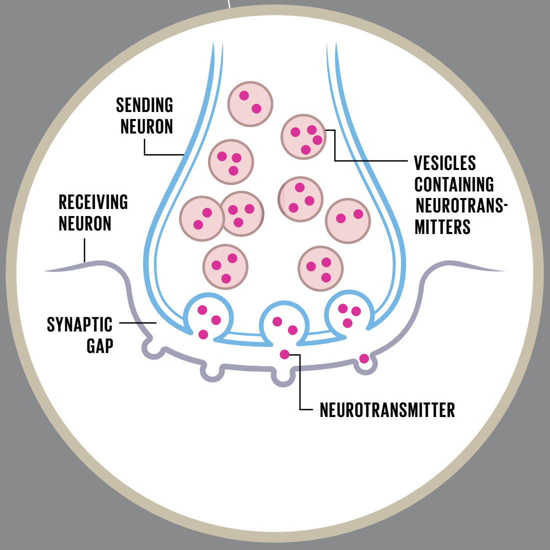

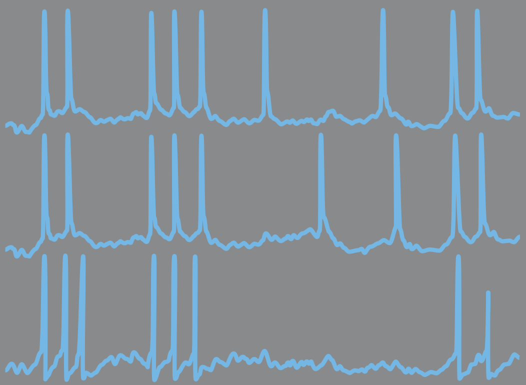

4 ACROSS THE GAP

In the majority of neurons, the action potential triggers the release of chemicals called neurotransmitters into the synapse. These molecules activate receptors on the connecting cells and influence the receiving neurons in either an excitatory or inhibitory fashion.

5 FIRING LINES

In the brain, millions of neurons are constantly activating, or firing, in complicated sequences. Neuroscientists are just beginning to decipher this neural code and to link firing patterns to sensations, actions, and feats of cognition.

ILLUSTRATIONS BY James Provost

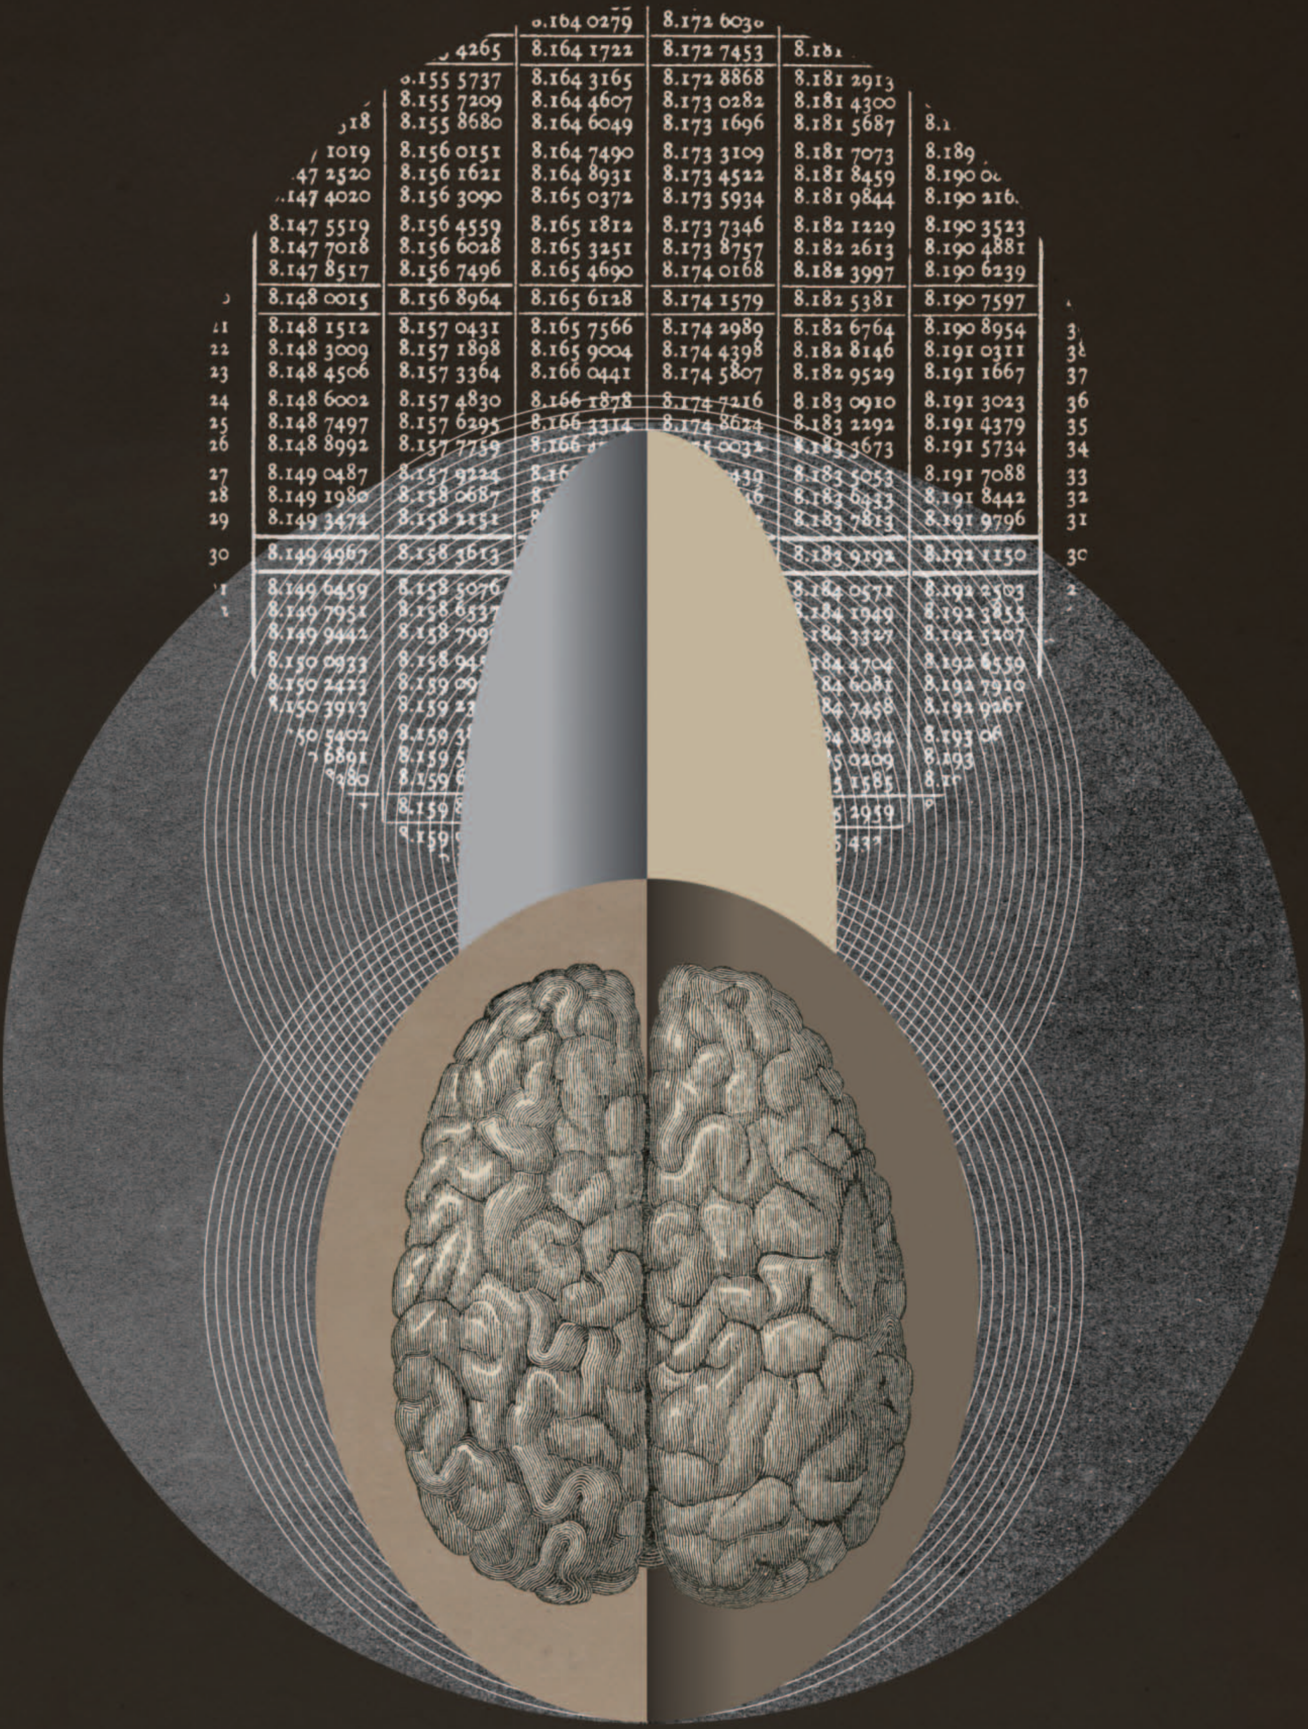
8.147 5519
8.147 7018
8.147 8517
8.148 0015
8.148 1512
8.148 3009
8.148 4506
8.148 6002
8.148 7497
8.148 8992
8.149 0487
8.149 1980
8.149 3474
8.149 4967
8.149 6459
8.149 7951
8.149 9442
8.150 0933
8.150 2423
8.155 5737
8.155 7209
8.155 8680
8.156 0151
8.156 1621
8.156 3090
8.156 4559
8.156 6028
8.156 7496
8.156 8964
8.157 0431
8.157 1898
8.157 3364
8.157 4830
8.164 1722
8.164 3165
8.164 4607
8.164 6049
8.164 7490
8.164 8931
8.165 0372
8.165 1812
8.165 3251
8.165 4690
8.165 6128
8.165 7566
8.165 9004
8.166 0441
8.172 7453
8.172 8868
8.173 0282
8.173 1696
8.173 3109
8.173 4522
8.173 5934
8.173 7346
8.173 8757
8.174 0168
8.174 1579
8.174 2989
8.174 4398
8.174 5807
8.181 2913
8.181 4300
8.181 5687
8.181 7073
8.181 8459
8.181 9844
8.182 1229
8.182 2613
8.182 3997
8.182 5381
8.182 6764
8.182 8146
8.182 9529
8.183 0910
8.183 2292
8.183 3673
8.190 3523
8.190 4881
8.190 6239
8.190 7597
8.190 8954
8.191 0311
8.191 1667
8.191 3023
8.191 4379
8.191 5734
8.191 7088
8.191 8442
8.191 9796
8.192 1150
8.192 2503
8.192 3855
8.192 5207
8.192 6559
8.192 7910
22
23
24
25
26
27
28
29
30
37
36
35
34
33
32
31

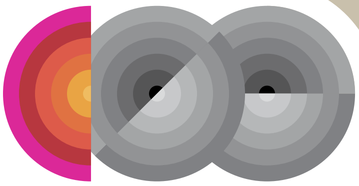

The Brain as Computer

THE BRAIN MAY BE BAD AT CRUNCHING NUMBERS, BUT IT'S A MARVEL OF COMPUTATIONAL EFFICIENCY

By **Karlheinz Meier**

PAINFUL EXERCISES IN BASIC ARITHMETIC ARE a vivid part of our elementary school memories. A multiplication like 3,752 × 6,901 carried out with just pencil and paper for assistance may well take up to a minute. Of course, today, with a cellphone always at hand, we can quickly check that the result of our little exercise is 25,892,552. Indeed, the processors in modern cellphones can together carry out more than 100 billion such operations per second. What's more, the chips consume just a few watts of power, making them vastly more efficient than our slow brains, which consume about 20 watts and need significantly more time to achieve the same result. ● Of course, the brain didn't evolve to perform arithmetic. So it does that rather badly. But it excels at processing a continuous stream of information from our surroundings. And it acts on that information—sometimes far more rapidly than we're aware of. No matter how much energy

ILLUSTRATION BY Chad Hagen

a conventional computer consumes, it will struggle with feats the brain finds easy, such as understanding language and running up a flight of stairs.

If we could create machines with the computational capabilities and energy efficiency of the brain, it would be a game changer. Robots would be able to move masterfully through the physical world and communicate with us in plain language. Large-scale systems could rapidly harvest large volumes of data from business, science, medicine, or government to detect novel patterns, discover causal relationships, or make predictions. Intelligent mobile applications like Siri or Cortana would rely less on the cloud. The same technology could also lead to low-power devices that can support our senses, deliver drugs, and emulate nerve signals to compensate for organ damage or paralysis.

But isn't it much too early for such a bold attempt? Isn't our knowledge of the brain far too limited to begin building technologies based on its operation? I believe that emulating even very basic features of neural circuits could give many commercially relevant applications a remarkable boost. How faithfully computers will have to mimic biological detail to approach the brain's level of performance remains an open question. But today's brain-inspired, or neuromorphic, systems will be important research tools for answering it.

* * *

A KEY FEATURE of conventional computers is the physical separation of memory, which stores data and instructions, from logic, which processes that information. The brain holds no such distinction. Computation and data storage are accomplished together locally in a vast network consisting of roughly 100 billion neural cells (neurons) and more than 100 trillion connections (synapses). Most of what the brain does is determined by those connections and by the manner in which each neuron responds to incoming signals from other neurons.

When we talk about the extraordinary capabilities of the human brain, we are usually referring to just the latest addition in the long evolutionary process that constructed it: the neocortex. This thin, highly folded layer forms the outer shell of our brains and carries out a diverse set of tasks that includes processing sensory inputs, motor control, memory, and learning. This great range of abilities is accomplished with a rather uniform structure: six horizontal layers and a million 500-micrometer-wide vertical columns all built from neurons, which integrate and distribute electrically coded information along tendrils that extend from them—the dendrites and axons.

Like all the cells in the human body, a neuron normally has an electric potential of about -70 millivolts between its interior and exterior. This membrane voltage changes when a neuron receives signals from other neurons connected to it. And if the membrane voltage rises to a critical threshold, it forms a voltage pulse, or spike, with a duration of a few milliseconds and a value of about 40 mV. This spike propagates along the neuron's axon until it reaches a synapse, the complex biochemical structure that connects the axon of one neuron to a dendrite of another. If the spike meets certain criteria, the synapse transforms it into another voltage pulse that travels down the branching dendrite structure of the receiving neuron and contributes either positively or negatively to its cell membrane voltage.

Connectivity is a crucial feature of the brain. The pyramidal cell, for example—a particularly important kind of cell in the human neocortex—contains about 30,000 synapses and so 30,000 inputs from other neurons. And the brain is constantly adapting. Neuron and synapse properties—and even the network structure itself—are always changing, driven mostly by sensory input and feedback from the environment.

General-purpose computers these days are digital rather than analog, but the brain is not as easy to categorize. Neurons accumulate electric charge just as capacitors in electronic circuits do. That is clearly an analog process. But the brain also uses spikes as units of information, and these are fundamentally binary: At any one place and time, there is either a spike or there is not. Electronically speaking, the brain is a mixed-signal system, with local analog computing and binary-spike communication. This mix of analog and digital helps the brain overcome transmission losses. Because the spike essentially has a value of either 0 or 1, it can travel a long distance without losing that basic information; it is also regenerated when it reaches the next neuron in the network.

Another crucial difference between brains and computers is that the brain accomplishes all its information processing without a central clock to synchronize it. Although we observe synchronization events—brain waves—they are self-organized, emergent products of neural networks. Interestingly, modern computing has started to adopt brainlike asynchronicity, to help speed up computation by performing operations in parallel. But the degree and the purpose of parallelism in the two systems are vastly different.

* * *

THE IDEA OF USING THE BRAIN as a model of computation has deep roots. The first efforts focused on a simple threshold neuron, which gives one value if the sum of weighted inputs is above a threshold and another if it is below. The biological realism of this scheme, which Warren McCulloch and Walter Pitts conceived in the 1940s, is very limited. Nonetheless, it was the first step toward adopting the concept of a firing neuron as an element for computation.

In 1957, Frank Rosenblatt proposed a variation of the threshold neuron called the perceptron. A network of integrating nodes (artificial neurons) is arranged in layers. The "visible" layers at the edges of the network interact with the outside world as inputs or outputs, and the "hidden" layers, which perform the bulk of the computation, sit in between.

Rosenblatt also introduced an essential feature found in the brain: inhibition. Instead of simply adding inputs together, the neurons in a perceptron network could also make negative contributions. This feature allows a neural network using only a single hidden layer to solve the XOR problem in logic, in which the output is true only if exactly one of the two binary inputs is true. This simple example shows that adding biological realism can add new computational capabilities. But which features of the brain are essential to what it can do, and which are just useless vestiges of evolution? Nobody knows.

We do know that some impressive computational feats can be accomplished without resorting to much biological realism. Deep-learning researchers have, for example, made great strides in using computers to analyze large volumes of data and pick out features in complicated images. Although the neural networks they build have more inputs and hidden layers than ever before, they are still based on the very simple neuron models. Their great capabilities reflect not their biological realism, but the scale of the networks they contain and the very powerful computers that are used to train them. But deep-learning networks are still a long way from the computational performance, energy efficiency, and learning capabilities of biological brains.

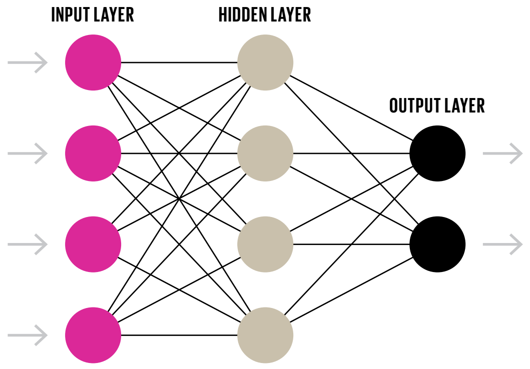

ARTIFICIAL NEURAL NETWORK: Today's machine learning applications are elaborations of this approach, in which interconnected nodes perform basic processing.

The big gap between the brain and today's computers is perhaps best underscored by looking at large-scale simulations of the brain. There have been several such efforts over the years, but they have all been severely limited by two factors: energy and simulation time. As an example, consider a simulation that Markus Diesmann and his colleagues conducted several years ago using nearly 83,000 processors on the K supercomputer in Japan. Simulating 1.73 billion neurons consumed 10 billion times as much energy as an equivalent size portion of the brain, even though it used very simplified models and did not perform any learning. And these simulations generally ran at less than a thousandth of the speed of biological real time.

Why so slow? The reason is that simulating the brain on a conventional computer requires billions of differential equations coupled together to describe the dynamics of cells and networks: analog processes like the movement of charges across a cell membrane. Computers that use Boolean logic–which trades energy for precision–and that separate memory and computing, appear to be very inefficient at truly emulating a brain.

These computer simulations can be a tool to help us understand the brain, by transferring the knowledge gained in the laboratory into simulations that we can experiment with and compare with real-world observations. But if we hope to go in the other direction and use the lessons of neuroscience to build powerful new computing systems, we have to rethink the way we design and build computers.

Copying brain operation in electronics may actually be more feasible than it seems at first glance. It turns out the energy cost of creating an electric potential in a synapse is about 10 femtojoules (10^{-15} joules). The gate of a metal-oxide-semiconductor (MOS) transistor that is considerably larger and more energy-hungry than those used in state-of-the-art CPUs requires just 0.5 fJ to charge. A synaptic transmission is therefore equivalent to the charging of at least 20 transistors. What's more, on the device level, biological and electronic circuits are not that different. So, in principle, we should be able to build structures like synapses and neurons from transistors and wire them up to make an artificial brain that wouldn't consume an egregious amount of energy.

* * *

THE NOTION OF BUILDING COMPUTERS by making transistors operate more like neurons began in the 1980s with Caltech professor Carver Mead. One of the core arguments behind what Mead came to call "neuromorphic" computing was that semiconductor devices can, when operated in a certain mode, follow the same physical rules as neurons

NEURONS IN SILICO

THE BRAINSCALES DESIGN evolved from this early prototype chip, dubbed "Spikey." There are four main components to the chip: neurons, which integrate signals from multiple synapses and fire when the voltage exceeds a threshold; line drivers that help condition the signals coming from neurons and send them to synapses; dense arrays of synapses that weigh those signals and have "plasticity," changing on short and long time scales; and digital circuitry to connect to the outside world and pass information such as chip configurations or data. The connections between neurons can be set externally or they can evolve internally according to plasticity mechanisms.

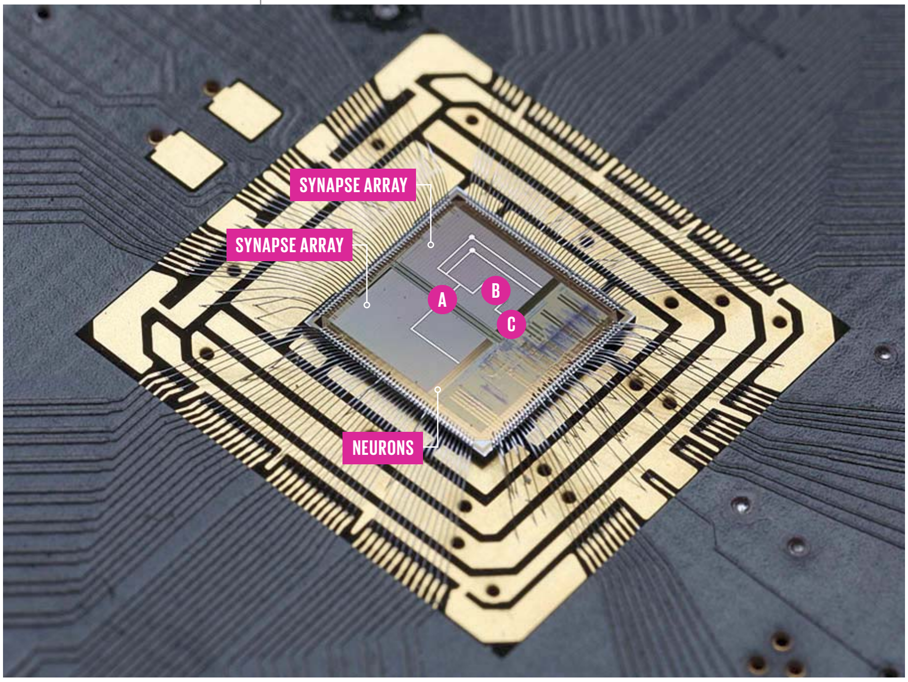

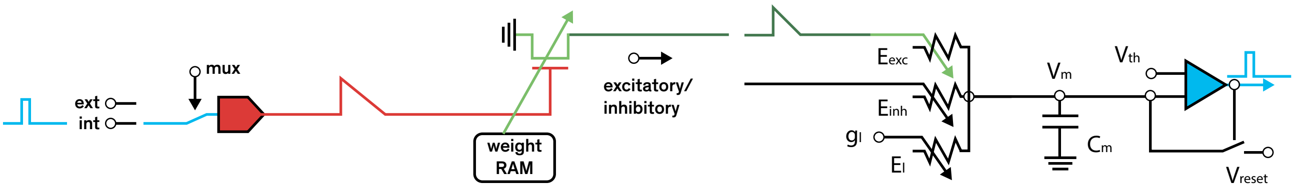

Ⓐ LINE DRIVER

An on-chip neuron or external circuitry produces a rectangular voltage pulse [blue]. The pulse is fed into a line driver [red icon], which converts the "spike" into a "postsynaptic potential"—a current pulse that falls off with a characteristic amount of time. The potential will be excitatory or inhibitory (with opposite voltage and direction of current flow) depending on settings determined by the user.

Ⓑ SYNAPSE

A transistor [green, open rectangle] fetches the weight that should be applied to the "postsynaptic potential" from a nearby memory cell. The pulse's shape, height, and other parameters are adjusted to better match signals created by biological synapses.

Ⓒ NEURON

Synapse pulses feed into each neuron, either contributing or subtracting from its membrane voltage (V_m). This membrane voltage, connected to ground through a capacitor, decays with time. If V_m exceeds the neuron's threshold voltage, a discriminator [blue triangle] will produce a rectangular voltage pulse and the neuron's membrane voltage will reset.

do and that this analog behavior could be used to compute with a high level of energy efficiency.

The Mead group also invented a neural communication framework in which spikes are coded only by their network addresses and the time at which they occur. This was groundbreaking work because it was the first to make time an essential feature of artificial neural networks. Time is a key factor in the brain: It takes time for signals to propagate and membranes to respond to changing conditions, and time determines the shape of postsynaptic potentials.

Several research groups quite active today, such as those of Giacomo Indiveri at ETH Zurich and Kwabena Boahen at Stanford, have followed Mead's approach and successfully implemented elements of biological cortical networks. The trick is to operate transistors below their turn-on threshold with extremely low currents, creating analog circuits that mimic neural behavior while consuming very little energy.

Further research in this direction may find applications in systems like brain-computer interfaces. But it would be a

HEIDELBERG UNIVERSITY

huge leap from there to circuitry that has anything like the network size, connectivity, and learning ability of a complete animal brain.

So, around 2005, three groups independently started to develop neuromorphic systems that deviate substantially from Mead's original approach, with the goal of creating large-scale systems with millions of neurons.

Closest to conventional computing is the SpiNNaker project, led by Steve Furber at the University of Manchester, in England. That group has designed a custom, fully digital chip that contains 18 ARM microprocessor cores operating with a clock speed of 200 megahertz–about a tenth the speed of modern CPUs. Although the ARM cores are classical computers, they simulate spikes, which are sent through specialized routers designed to communicate asynchronously, just as the brain does. The current implementation, part of the European Union's Human Brain Project, was completed in 2016 and features 500,000 ARM cores. Depending on the complexity of the neuron model, each core can simulate as many as 1,000 neurons.

The TrueNorth chip, developed by Dharmendra Modha and his colleagues at the IBM Research laboratory in Almaden, Calif., abandons the use of microprocessors as computational units and is a truly neuromorphic computing system, with computation and memory intertwined. TrueNorth is still a fully digital system, but it is based on custom neural circuits that implement a very specific neuron model. The chip features 5.4 billion transistors, built with 28-nanometer Samsung CMOS technology. These transistors are used to implement 1 million neuron circuits and 256 million simple (1-bit) synapses on a single chip.

The furthest away from conventional computing and closest to the biological brain, I would argue, is the BrainScaleS system, which my colleagues and I have developed at Heidelberg University, in Germany, for the Human Brain Project. BrainScaleS is a mixed-signal implementation. That is, it combines neurons and synapses made from silicon transistors that operate as analog devices with digital communication. The full-size system is built from 20 uncut 8-inch silicon wafers to create a total of 4 million neurons and 1 billion synapses.

The system can replicate eight different firing modes of biological neurons, derived in close collaboration with neuroscientists. Unlike with the analog approach pioneered by Mead, BrainScaleS operates in an accelerated mode, running about 10,000 times as fast as real time. This makes it especially suited to study learning and development.

Learning will likely be a critical component of neuromorphic systems going forward. At the moment, brain-inspired chips as well as the neural networks implemented on conventional computers are trained elsewhere by more powerful computers. But if we want to use neuromorphic systems in real-world applications, to, say, power robots that will work alongside us, they will have to be able to learn and adapt on the fly.

In our second-generation BrainScaleS system, my colleagues and I implemented learning capabilities by building on-chip "plasticity processors," which are used to alter neuron and synapse parameters when needed. This capability also lets us fine-tune parameters to compensate for differences in size and electric properties from device to device, much as the brain compensates for variation.

The three large-scale systems I've described are complementary approaches. SpiNNaker is highly configurable and so can be used to test a variety of different neuron models, TrueNorth has very high integration density, and BrainScaleS is designed for continuous operation of learning and development. Finding the right way to evaluate the performance of such systems is an ongoing effort. But early results indicate promise. The IBM TrueNorth group, for example, recently estimated that a synaptic transmission in its system costs 26 picojoules. Although this is about a thousand times the energy of the same action in a biological system, it is approaching 1/100,000 of the energy that would be consumed by a simulation carried out on a conventional general-purpose machine.

We are still in the early days of figuring out what these sorts of systems can do and how they can be applied to real-world applications. At the same time, we must find ways to combine many neuromorphic chips into larger networks with improved learning capabilities, while simultaneously driving down energy consumption. One challenge is simply connectivity: The brain is three-dimensional, and we build circuits in two. Three-dimensional integration of circuits, a field that is actively being explored, could help with this.

Another enabler will be non-CMOS devices, such as memristors or phase-change RAM. Today, the weight values that govern how artificial synapses respond to incoming signals are stored in conventional digital memory, which dominates the silicon resources required to build a network. But other forms of memory could let us shrink those cells down from micrometer to nanometer scales. As in today's systems, a key challenge there will be how to handle variations between individual devices. The calibration strategies developed with BrainScaleS could help.

We are just getting started on the road toward usable and useful neuromorphic systems. But the effort is well worthwhile. If we succeed, we won't just be able to build powerful computing systems; we may even gain insights about our own brains. ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/brainascomputer0617


NEOCORTEX FOR NEOPHYTES
Neocortex for Neophytes
Information for the Neuro Novice

What Intelligent Machines Need to Learn From the Neocortex

MACHINES WON'T BECOME INTELLIGENT UNLESS THEY INCORPORATE CERTAIN FEATURES OF THE HUMAN BRAIN. HERE ARE THREE OF THEM

By **Jeff Hawkins**

COMPUTERS HAVE TRANSFORMED work and play, transportation and medicine, entertainment and sports. Yet for all their power, these machines still cannot perform simple tasks that a child can do, such as navigating an unknown room or using a pencil. ● The solution is finally coming within reach. It will emerge from the intersection of two major pursuits: the reverse engineering of the brain and the burgeoning field of artificial intelligence. Over the next 20 years, these two pursuits will combine to usher in a new epoch of intelligent machines. ● Why do we need to know how the brain works to build intelligent machines? Although machine-learning techniques such as deep neural networks have recently made impressive gains, they are still a world away from being intelligent, from being able to understand and act in the world the way that we do. The only example of intelligence, of the ability to learn from the world, to plan and to execute, is the

PHOTOGRAPH BY Dan Saelinger

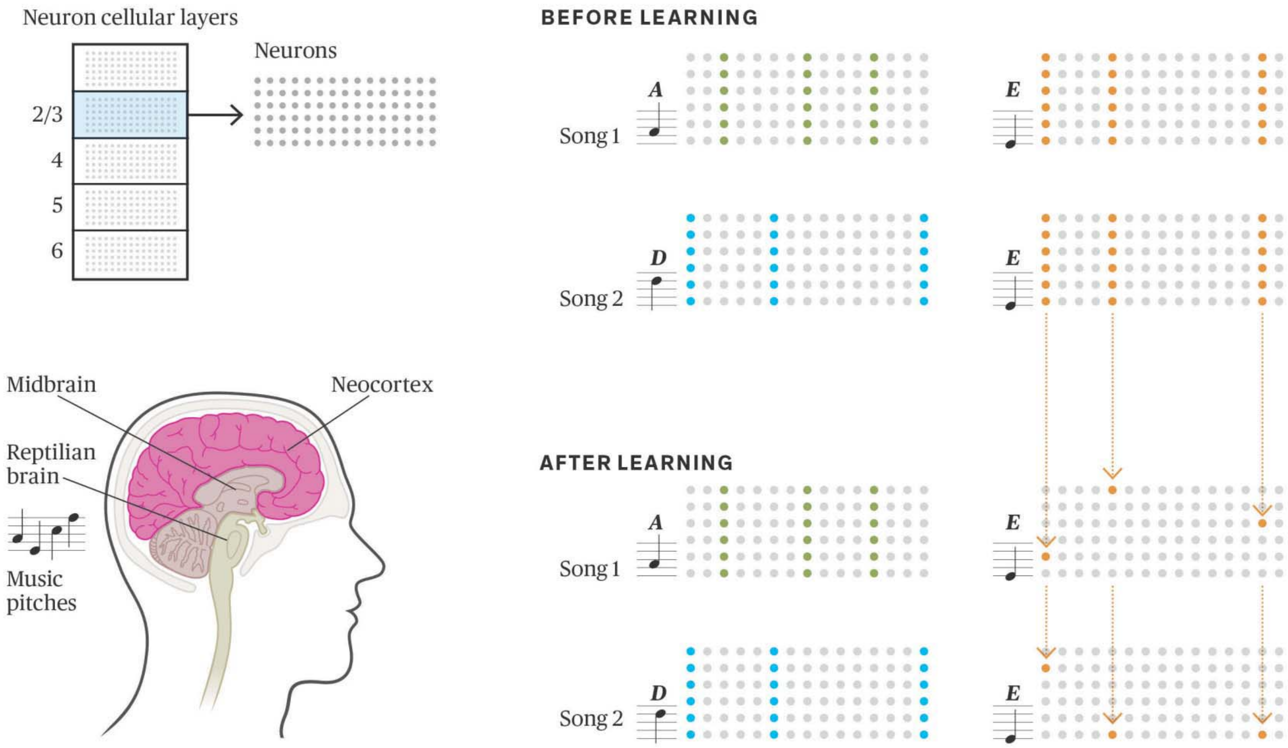

THE SOUND OF MUSIC, like all sensory experience, is created in the brain by patterns of neurons firing in the brain's neocortex [illustration, lower left], which has multiple layers [upper left]. In this diagram a layer is shown schematically; in reality, each layer has thousands of neurons. When a song is heard for the first time, each musical note triggers a specific pattern of neuronal firing, as shown in these diagrams, in which a colored dot indicates a firing neuron. The top two lines show the firing patterns associated with a listener hearing, for the first time, the first four notes of two different songs. In these two songs, the second note of the first song is the same as the second note of the second song. Similarly, the third notes are the same. After the listener recognizes a melody ["After Learning"], only one neuron in the column fires in response to the notes [arrows], as shown in the bottom two lines. That's because these notes are each unique representations—for example, the note E after an A in a specific song. More details are available at http://numenta.com/neuron-paper.

brain. Therefore, we must understand the principles underlying human intelligence and use them to guide us in the development of truly intelligent machines.

* * *

AT MY COMPANY, Numenta, in Redwood, City, Calif., we study the neocortex, the brain's largest component and the one most responsible for intelligence. Our goal is to understand how it works and to identify the underlying principles of human cognition. In recent years, we have made significant strides in our work, and we have identified several features of biological intelligence that we believe will need to be incorporated into future thinking machines.

To understand these principles, we must start with some basic biology. The human brain is similar to a reptile's brain. Each has a spinal cord, which controls reflex behaviors; a brain stem, which controls autonomic behaviors such as breathing and heart rate; and a midbrain, which controls emotions and basic behaviors. But humans, indeed all mammals, have something reptiles don't: a neocortex.

The neocortex is a deeply folded sheet some 2 millimeters thick that, if laid out flat, would be about as big as a large dinner napkin. In humans, it takes up about 75 percent of the brain's volume. This is the part that makes us smart.

At birth, the neocortex knows almost nothing; it learns through experience. Everything we learn about the world—driving a car, operating a coffee machine, and the thousands

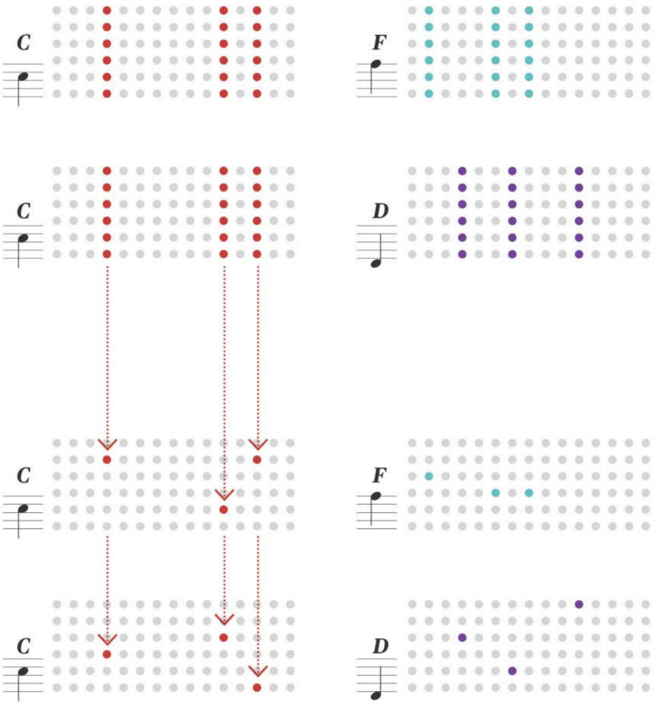

of other things we interact with every day–is stored in the neocortex. It learns what these objects are, where they are in the world, and how they behave. The neocortex also generates motor commands, so when you make a meal or write software it is the neocortex controlling these behaviors. Language, too, is created and understood by the neocortex.

The neocortex, like all of the brain and nervous system, is made up of cells called neurons. Thus, to understand how the brain works, you need to start with the neuron. Your neocortex has about 30 billion of them. A typical neuron has a single tail-like axon and several treelike extensions called dendrites. If you think of the neuron as a kind of signaling system, the axon is the transmitter and the dendrites are the receivers. Along the branches of the dendrites lie some 5,000 to 10,000 synapses, each of which connects to counterparts on thousands of other neurons. There are thus more than 100 trillion synaptic connections.

Your experience of the world around you–recognizing a friend's face, enjoying a piece of music, holding a bar of soap in your hand–is the result of input from your eyes, ears, and other sensory organs traveling to your neocortex and causing groups of neurons to fire. When a neuron fires, an electrochemical spike travels down the neuron's axon and crosses synapses to other neurons. If a receiving neuron gets enough input, it might then fire in response and activate other neurons. Of the 30 billion neurons in the neocortex, 1 or 2 percent are firing at any given instant, which means that many millions of neurons will be active at any point in time. The set of active neurons changes as you move and interact with the world. Your perception of the world, what you might consider your conscious experience, is determined by the constantly changing pattern of active neurons.

The neocortex stores these patterns primarily by forming new synapses. This storage enables you to recognize faces and places when you see them again, and also recall them from your memory. For example, when you think of your friend's face, a pattern of neural firing occurs in the neocortex that is similar to the one that occurs when you are actually seeing your friend's face.

Remarkably, the neocortex is both complex and simple at the same time. It is complex because it is divided into dozens of regions, each responsible for different cognitive functions. Within each region there are multiple layers of neurons, as well as dozens of neuron types, and the neurons are connected in intricate patterns.

The neocortex is also simple because the details in every region are nearly identical. Through evolution, a single algorithm developed that can be applied to all the things a neocortex does. The existence of such a universal algorithm is exciting because if we can figure out what that algorithm is, we can get at the heart of what it means to be intelligent, and incorporate that knowledge into future machines.

But isn't that what AI is already doing? Isn't most of AI built on "neural networks" similar to those in the brain? Not really. While it is true that today's AI techniques reference neuroscience, they use an overly simplified neuron model, one that omits essential features of real neurons, and they are connected in ways that do not reflect the reality of our brain's complex architecture. These differences are many, and they matter. They are why AI today may be good at labeling images or recognizing spoken words but is not able to reason, plan, and act in creative ways.

Our recent advances in understanding how the neocortex works give us insights into how future thinking machines will work. I am going to describe three aspects of biological intelligence that are essential, but largely missing from today's AI. They are learning by rewiring, sparse representations, and embodiment, which refers to the use of movement to learn about the world.

CHART SOURCE: JEFF HAWKINS; ILLUSTRATION: MARK MONTGOMERY

HEAD ILLUSTRATION BY **James Provost**

◆

LEARNING BY REWIRING: Brains exhibit some remarkable learning properties. First, we learn quickly. A few glances or a few touches with the fingers are often sufficient to learn something new. Second, learning is incremental. We can learn something new without retraining the entire brain or forgetting what we learned before. Third, brains learn continuously. As we move around the world, planning and acting, we never stop learning. Fast, incremental, and continuous learning are essential ingredients that enable intelligent systems to adapt to a changing world. The neuron is responsible for learning, and the complexities of real neurons are what make it a powerful learning machine.

In recent years, neuroscientists have learned some remarkable things about the dendrite. One is that each of its branches acts as a set of pattern detectors. It turns out that just 15 to 20 active synapses on a branch are sufficient to recognize a pattern of activity in a large population of neurons. Therefore, a single neuron can recognize hundreds of distinct patterns. Some of these recognized patterns cause the neuron to become active, but others change the internal state of the cell and act as a prediction of future activity.

Neuroscientists used to believe that learning occurred solely by modifying the effectiveness of existing synapses so that when an input arrived at a synapse it would either be more likely or less likely to make the cell fire. However, we now know that most learning results from growing new synapses between cells–by “rewiring” the brain. Up to 40 percent of the synapses on a neuron are replaced with new ones every day. New synapses result in new patterns of connections among neurons, and therefore new memories. Because the branches of a dendrite are mostly independent, when a neuron learns to recognize a new pattern on one of its dendrites, it doesn’t interfere with what the neuron has already learned on other dendrites.

This is why we can learn new things without interfering with old memories and why we don’t have to retrain the brain every time we learn something new. Today’s neural networks don’t have these properties.

Intelligent machines don’t have to model all the complexity of biological neurons, but the capabilities enabled by dendrites and learning by rewiring are essential. These capabilities will need to be in future AI systems.

◆

SPARSE REPRESENTATIONS: Brains and computers represent information quite differently. In a computer’s memory, all combinations of 1s and 0s are potentially valid, so if you change one bit it will typically result in an entirely different meaning, in much the same way that changing the letter *i* to *a* in the word *fire* results in an unrelated word, *fare*. Such a representation is therefore brittle.

Brains, on the other hand, use what’s called sparse distributed representations, or SDRs. They’re called sparse because relatively few neurons are fully active at any given time. Which neurons are active changes moment to moment as you move and think, but the percentage is always small. If we think of each neuron as a bit, then to represent a piece of information the brain uses thousands of bits (many more than the 8 to 64 used in computers), but only a small percentage of the bits are 1 at any time; the rest are 0.

Let’s say you want to represent the concept of “cat” using an SDR. You might use 10,000 neu-

THE FUTURE according to...

Martine Rothblatt

Founder, Terasem Movement, which promotes transhumanism

When will we have computers as capable as the brain?
In the 21st century.

How will brainlike computers change the world?
People will no longer die, because the computerized version of themselves will insist they are still alive even if their flesh body has expired. They will be sad, and they’ll be hoping for regenerated bodies, but they will feel like a backup of their original flesh self.

Do you have any qualms about a future in which computers have human-level (or greater) intelligence?
I don’t, because I’m confident the computers will be overwhelmingly friendly since they will be selected for in a Darwinian environment that consists of humanity. There is no market for [a] bad robot, no more than there is for a bad car or plane. Of course, there is the inevitability of a DIY bad human-level computer, but that gives me even more reason to welcome human-level cyberintelligence, because just like it takes a [thief to catch a thief], it will take a human-friendly smart computer to catch an antihuman smart computer.

PHOTO-ILLUSTRATIONS BY Gluekit

rons of which 100 are active. Each of the active neurons represents some aspect of a cat, such as "pet," or "furry," or "clawed." If a few neurons die, or a few extra neurons become active, the new SDR will still be a good representation of "cat" because most of the active neurons are still the same. SDRs are thus not brittle but inherently robust to errors and noise. When we build silicon versions of the brain, they will be intrinsically fault tolerant.

There are two properties of SDRs I want to mention. One, the overlap property, makes it easy to see how two things are similar or different in meaning. Imagine you have one SDR representing "cat" and another representing "bird." Both the "cat" and "bird" SDR would have the same active neurons representing "pet" and "clawed," but they wouldn't share the neuron for "furry." This example is simplified, but the overlap property is important because it makes it immediately clear to the brain how the two objects are similar or different. This property confers the power to generalize, a capability lacking in computers.

The second, the union property, allows the brain to represent multiple ideas simultaneously. Imagine I see an animal moving in the bushes, but I got only a glimpse, so I can't be sure of what I saw. It might be a cat, a dog, or a monkey. Because SDRs are sparse, a population of neurons can activate all three SDRs at the same time and not get confused, because the SDRs will not interfere with one another. The ability of neurons to constantly form unions of SDRs makes them very good at handling uncertainty.

Such properties of SDRs are fundamental to understanding, thinking, and planning in the brain. We can't build intelligent machines without embracing SDRs.

* * *

EMBODIMENT: The neocortex receives input from the sensory organs. Every time we move our eyes, limbs, or body, the sensory inputs change. This constantly changing input is the primary mechanism the brain uses to learn about the world. Imagine I present you with an object you have never seen before. For the sake of discussion, let's say it's a stapler. How would you learn about the new object? You might walk around the stapler, looking at it from different angles. You might pick it up, run your fingers over it, and rotate it in your hands. You then might push and pull on it to see how it behaves. Through this interactive process, you learn the shape of the stapler, what it feels like, what it looks like, and how it behaves. You make a movement, see how the inputs change, make another movement, see how the inputs change again, and so on. Learning through movement is the brain's primary means for learning. It will be a central component of all truly intelligent systems.

This is not to say that an intelligent machine needs a physical body, only that it can change what it senses by moving. For example, a virtual AI machine could "move" through the Web by following links and opening files. It could learn the structure of a virtual world through virtual movements, analogous to what we do when walking through a building.

This brings us to an important discovery we made at Numenta last year. In the neocortex, sensory input is processed in a hierarchy of regions. As sensory input passes from one level of the hierarchy to another, more complex features are extracted, until at some point an object can be recognized. Deep-learning networks also use hierarchies, but they often require 100 levels of processing to recognize an image, whereas the neocortex achieves the same result with just four levels. Deep-learning networks also require millions of training patterns, while the neocortex can learn new objects with just a few movements and sensations. The brain is doing something fundamentally different than a typical artificial neural network, but what?

Hermann von Helmholtz, the 19th-century German scientist, was one of the first people to suggest an answer. He observed that, although our eyes

| CONTINUED ON PAGE 68

THE FUTURE according to...

Ruchir Puri

Chief architect, IBM Watson

When will we have computers as capable as the brain?

A human brain is fundamentally different than being a champion chess, "Jeopardy!," or Go player. It is something that entails essential traits like caring, empathy, sharing, ingenuity, and innovation. These human brain traits might prove to be elusive to machines for a long time, not to mention the incredible form factor and energy efficiency with which [the] human brain operates and learns. Although AI's impact on society will accelerate further...to reach new heights in solving problems which are currently grand challenges for human society, it will be a while before we will be able to holistically answer [that] question.

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/intelligentmachines0617

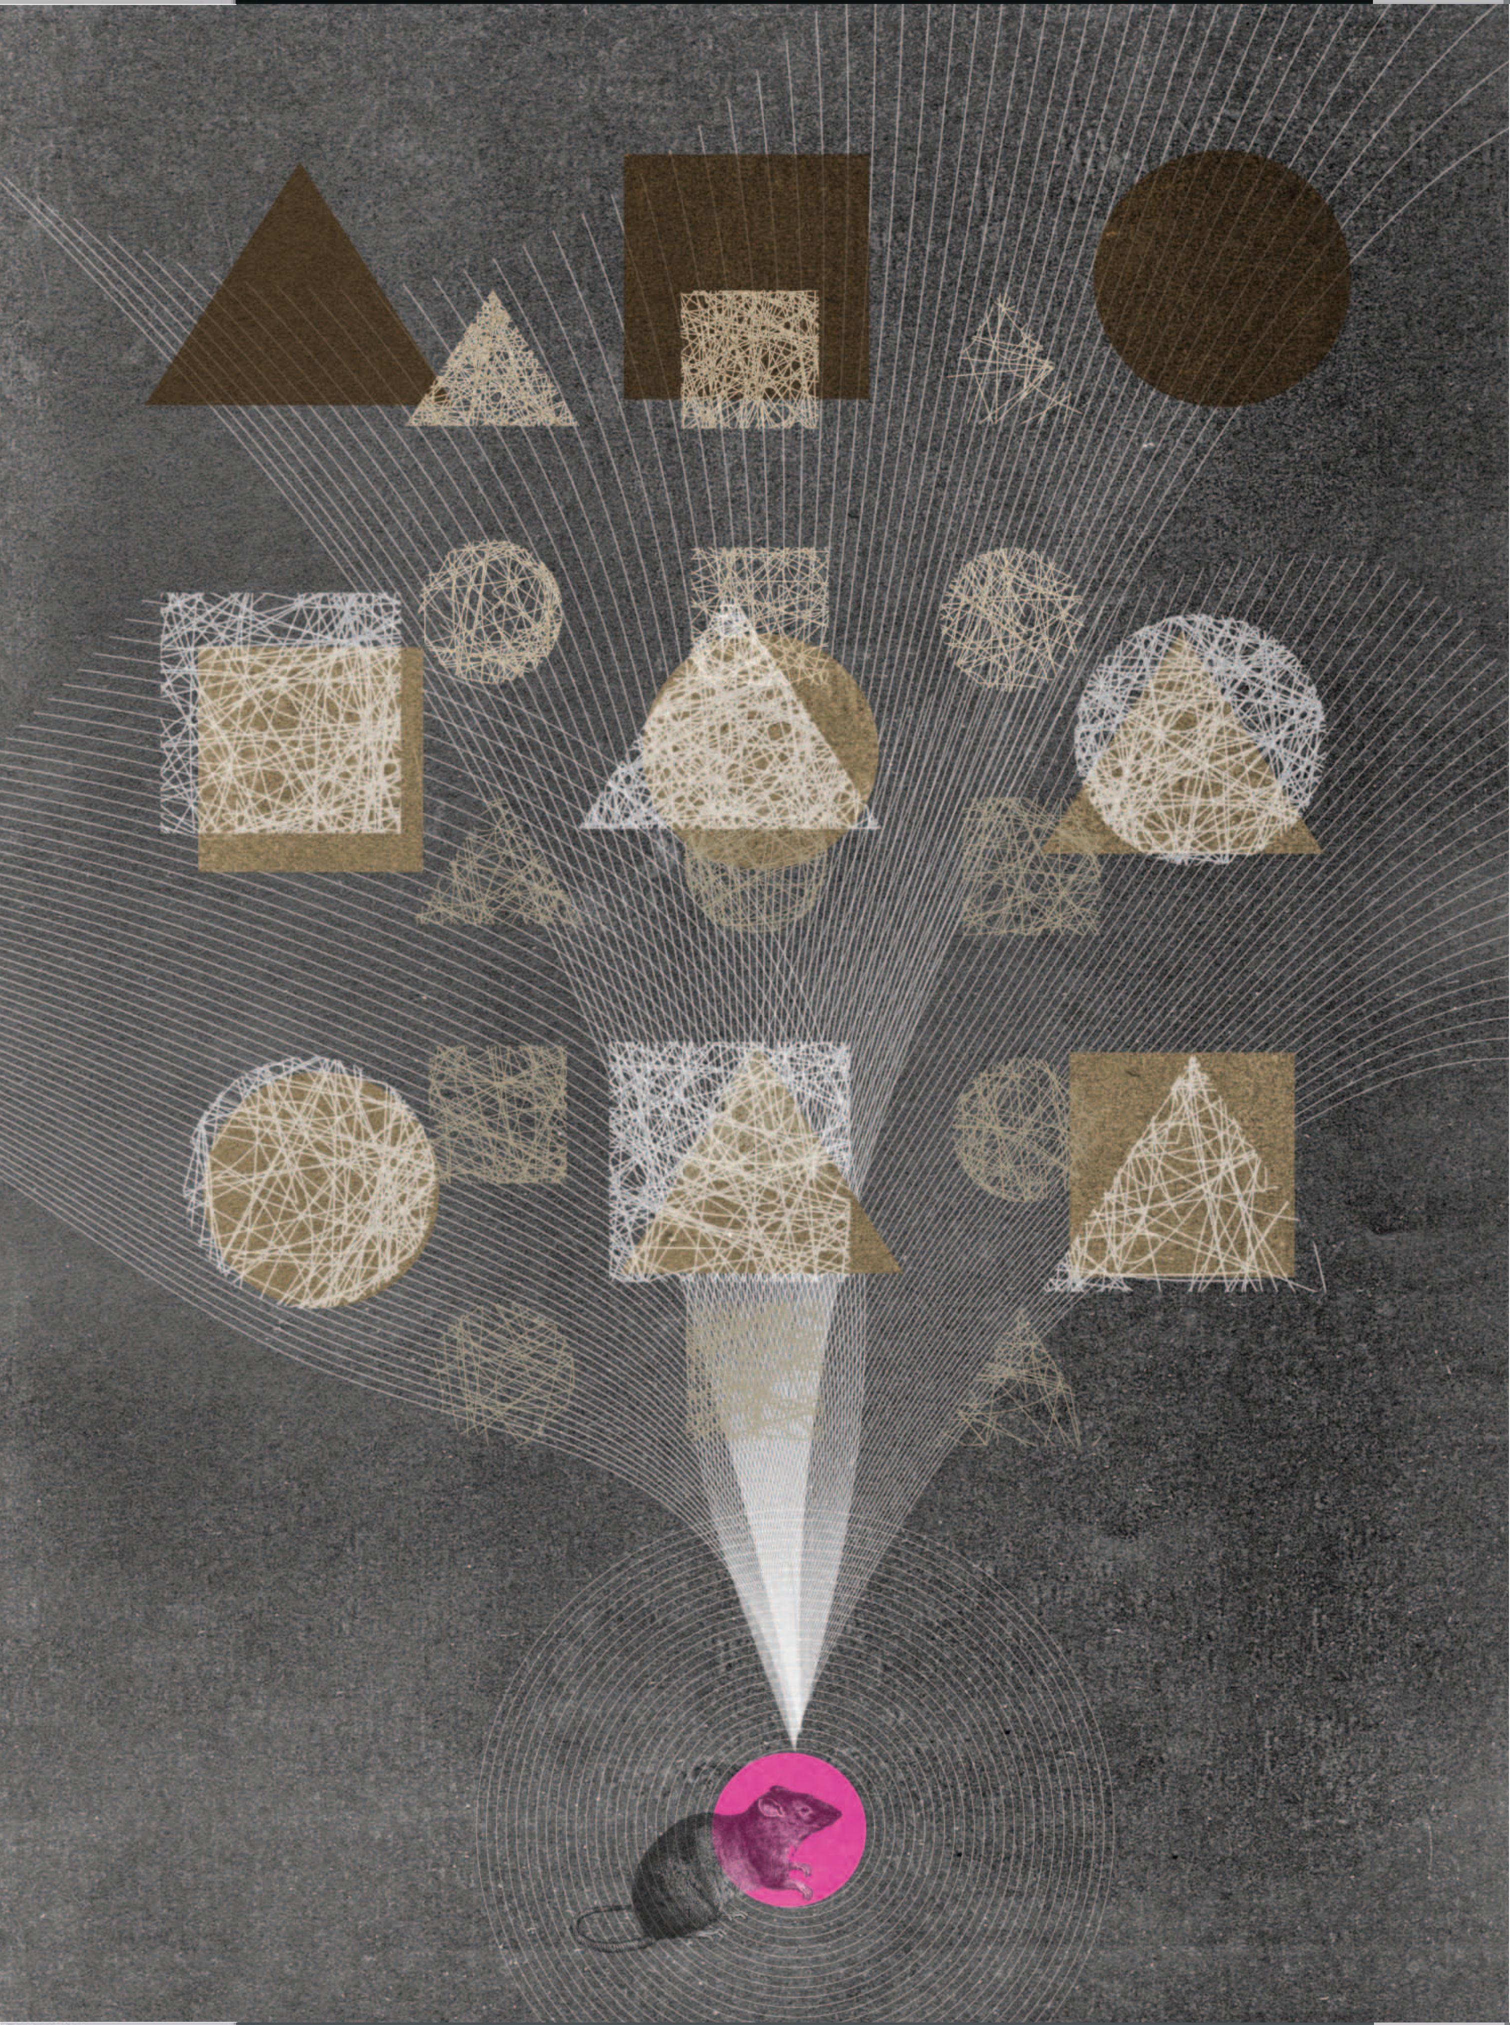

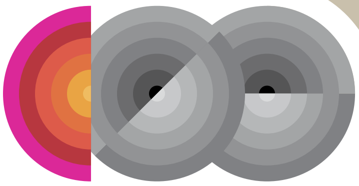

From Animal Intelligence to Artificial Intelligence

RESEARCHERS ARE TRYING TO FIGURE OUT WHAT MAKES THE BRAIN SPECIAL, SO THEY CAN BESTOW THAT SPECIALNESS ON COMPUTERS

By **Eliza Strickland**

TODAY'S ARTIFICIAL INTELLIGENCE systems can destroy human champions at sophisticated games like chess, Go, and Texas Hold 'em. In flight simulators, they can shoot down top fighter pilots. They're surpassing human doctors with more precise surgical stitching and more accurate cancer diagnoses. But there are some situations when a 3-year-old can easily defeat the fanciest AI in the world: when the contest involves a type of learning so routine that humans don't even realize they're doing it. ● That last thought occurred to David Cox–Harvard neuroscientist, AI expert, and proud father of a 3-year-old–when his daughter spotted a long-legged skeleton at a natural history museum, pointed at it, and said, "Camel!" Her only other encounter with a camel had been a few months earlier, when he showed her a cartoon camel in a picture book. ● AI researchers call this ability to identify an object based on a single

ILLUSTRATION BY Chad Hagen

example "one-shot learning," and they're deeply envious of toddlers' facility with it. Today's AI systems acquire their smarts through a very different process. Typically, in an autonomous training method called deep learning, a program is given masses of data from which to draw conclusions. To train an AI camel detector, the system would ingest thousands of images of camels–cartoons, anatomical diagrams, photos of the one-humped and two-humped varieties–all of which would be labeled "camel." The AI would also take in thousands of other images labeled "not camel." Once it had chewed through all that data to determine the animal's distinguishing features, it would be a whiz-bang camel detector. But Cox's daughter would have long since moved on to giraffes and platypuses.

Cox mentions his daughter by way of explaining a U.S. government program called Machine Intelligence from Cortical Networks (Microns). Its ambitious goal: to reverse engineer human intelligence so that computer scientists can build better artificial intelligence. First, neuroscientists are tasked with discovering the computational strategies at work in the brain's squishy gray matter; then the data team will translate those strategies into algorithms. One of the big challenges for the resulting AIs will be one-shot learning. "Humans have an amazing ability to make inferences and generalize," Cox says, "and that's what we're trying to capture."

The five-year program, funded to the tune of US $100 million by the Intelligence Advanced Research Projects Agency (IARPA), keeps a tight focus on the visual cortex, the part of the brain where much visual-information processing occurs. Working with mice and rats, three Microns teams aim to map the layout of neurons inside 1 cubic millimeter of brain tissue. That may not sound like much, but that tiny cube contains about 50,000 neurons connected to one another at about 500 million junctures called synapses. The researchers hope that a clear view of all those connections will allow them to discover the neural "circuits" that are activated when the visual cortex is hard at work. The project requires specialized brain imaging that shows individual neurons with nanometer-level resolution, which has never before been attempted for a brain chunk of this size.

Although each Microns team involves multiple institutions, most of the participants of the team led by Cox, an assistant professor of molecular and cellular biology and of computer science at Harvard University, work in a single building on the Harvard campus. A tour through the halls reveals rodents busy with assignments in a rat "video arcade"; a machine that carves up brains as if it were the world's most precise deli slicer; and some of the fastest and most powerful microscopes on the planet. With their equipment running full out and with huge amounts of human effort, Cox thinks they just might crack the code of that daunting cubic millimeter.

● ■ ●

TRY TO WRAP YOUR MIND around the sheer power of the human mind. To process information about the world and to keep your body running, electric pulses flash through 86 billion neurons packed into spongy folds of tissue inside your cranium. Each neuron has a long axon that winds through the tissue and allows it to link to thousands of other neurons, mak-

TWO-PHOTON EXCITATION MICROSCOPY

A powerful infrared laser scans the brain tissue of a living animal while it's engaged in a task. When two photons simultaneously hit an active neuron, they cause a fluorescent tag to emit a photon of light at another wavelength. The microscope records these flashes of light. "You can watch a rat having a thought," says David Cox.

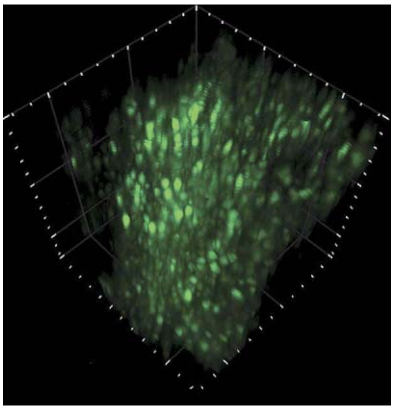

X-RAY TOMOGRAPHY

At Argonne National Laboratory's Advanced Photon Source, a particle accelerator slams electrons into a metal filament to produce extremely bright X-rays, which are focused on a tiny chunk of extracted brain tissue. X-ray images from many angles are combined to create a 3D image showing every neuron inside the chunk.

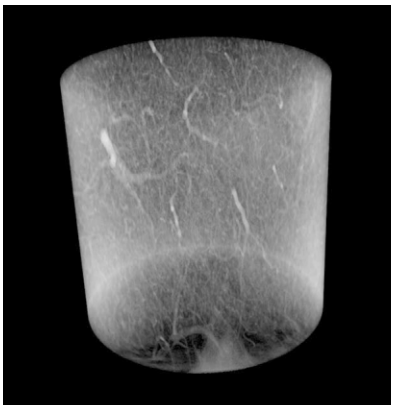

ing for trillions of connections. Patterns of electric pulses correlate with everything a human being experiences: wiggling a finger, digesting lunch, falling in love, recognizing a camel.

Computer scientists have tried to emulate the brain since the 1940s, when they first devised software structures called artificial neural networks. Most of today's fanciest AIs use some modern form of this architecture: There are deep neural networks, convolutional neural networks, recurrent neural networks, and so on. Loosely inspired by the brain's structure, these networks consist of many computing nodes called artificial neurons, which perform small discrete tasks and connect to each other in ways that allow the overall systems to accomplish impressive feats.

Neural networks haven't been able to copy the anatomical brain more closely because science still lacks basic information about neural circuitry. R. Jacob Vogelstein, the Microns project manager at IARPA, says researchers have typically worked on either the micro or macro scale. "The tools we used either involved poking individual neurons or aggregating signals across large swaths of the brain," he says. "The big gap is understanding operations on a circuit level–how thousands of neurons work together to process information."

The situation changed with recent technological advances that enable neuroscientists to make "connectome" maps revealing the multitude of connections among neurons. But Microns seeks more than a static circuitry diagram. The teams must show how those circuits activate when a rodent sees, learns, and remembers. "It's very similar to how you'd try to reverse engineer an integrated circuit," Vogelstein says. "You could stare at the chip in extreme detail, but you won't really know what it's meant to do unless you see the circuit in operation."

For IARPA, the real payoff will come if researchers can trace the pattern of neurons involved in a cognitive task and translate that pattern into a more brainlike architecture for artificial neural networks. "Hopefully, the brain's computational strategies are representable in mathematical and algorithmic terms," Vogelstein says. The government's big bet is that brainlike AI systems will be more adept than their predecessors at solving real-world problems. After all, it's a noble quest to understand the brain, but intelligence agencies want AIs that can quickly learn to recognize not just a camel but also a half-obscured face in a grainy security-camera video.

● ■ ●

THE VIDEO ARCADE for Cox's rats is a small room where black boxes the size of microwaves are stacked four high. Inside each box stands a rat facing a computer screen, with two nozzles directly in front of its nose.

In the current experiment, the rat tries to master a sophisticated visual task. The screen displays three-dimensional computer-generated objects–nothing recognizable from the outside world, just lumpy abstract shapes. When the rat sees object A, it must lick the nozzle on the left to get a drop of sweet juice; when it sees object B, the juice will be in the right nozzle. But the objects are presented in various orientations, so the rat has to mentally rotate each shape on display and decide if it matches A or B.

Interspersed with training sessions are imaging sessions, for which the rats are taken down the hall to another lab where a bulky microscope is draped in black cloth, looking like an old-fashioned photographer's setup. Here, the team uses a two-photon excitation microscope to examine the animal's visual cortex while it's looking at a screen displaying the now-familiar objects A and B, again in various orientations. The microscope records flashes of fluorescence when its laser hits active neurons, and the 3D video shows patterns that resemble green fireflies winking on and off in a summer night. Cox is keen to see how those patterns change as the animal becomes expert at its task.

The microscope's resolution isn't fine enough to show the axons that connect the neurons to one another. Without that information, the researchers can't determine how one neuron triggers the next to create an information-processing circuit. For that next step, the animal must be killed, and its brain subjected to much closer study.

The researchers excise a tiny cube of the visual cortex, which goes via FedEx to Argonne National Laboratory, in

SCANNING ELECTRON MICROSCOPY

Jeff Lichtman's scanning electron microscope fires 61 beams of electrons at thin slices of brain tissue. By measuring how the electrons scatter, this technique produces images at 4-nanometer resolution, showing the axons that connect neurons in each slice. A computer program traces axons from one slice to the next.

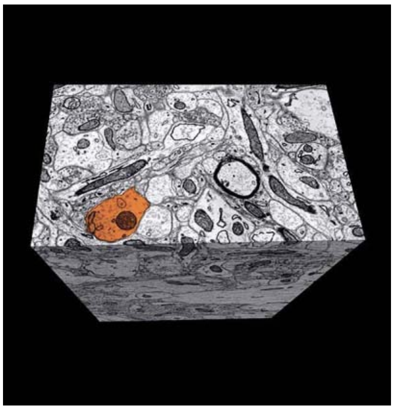

FROM LEFT: VAZIRI LAB; RAFAEL VESCOVI AND NARAYANAN KASTHURI/ARGONNE NATIONAL LABORATORY; DANIEL BERGER/LICHTMAN LAB/HARVARD UNIVERSITY

Illinois. There, a particle accelerator uses powerful X-ray radiation to make a 3D map showing the individual neurons, other types of brain cells, and blood vessels. This map also doesn't reveal the connecting axons inside the cube, but it helps later, when researchers compare the two-photon microscopy images with images produced with electron microscopes. "The X-ray is like a Rosetta stone," Cox says.

Then the brain nugget comes back to the Harvard lab of Jeff Lichtman, a professor of molecular and cellular biology and a leading expert on the brain's connectome. Lichtman's team takes that 1 mm^3 of brain and uses the machine that resembles a deli slicer to carve 33,000 slices, each only 30 nanometers thick. These gossamer sheets are automatically collected on strips of tape and arranged on silicon wafers. Next the researchers deploy one of the world's fastest scanning electron microscopes, which slings 61 beams of electrons at each brain sample and measures how the electrons scatter. The refrigerator-size machine runs around the clock, producing images of each slice with 4-nm resolution.

Each image resembles a cross section of a cube of densely packed spaghetti. Image-processing software arranges the slices in order and traces each strand of spaghetti from one slice to the next, delineating the full length of each neuron's axon along with its thousands of connections to other neurons. But the software sometimes loses track of strands or gets one confused with another. Humans are better at this task than computers, Cox says. "Unfortunately, there aren't enough humans on earth to trace this much data." Software engineers at Harvard and MIT are working on the tracing problem, which they must solve to make an accurate wiring diagram of the brain.

Overlaying that diagram with activity maps from the two-photon microscope should reveal the brain's computational structures. For example, it should show which neurons form a circuit that lights up when a rat sees an odd lumpy object, mentally flips it upside down, and decides that it's a match for object A.

Another big challenge for Cox's team is speed. In the program's first phase, which ended in May, each team had to show off results from a chunk of brain tissue that measured 100 micrometers cubed. For that smaller chunk, Cox's team had the electron-microscopy and image-reconstruction step down to two weeks. Now, in phase two, the teams need to be able to process the same size chunk in a few hours. Scaling up from 100 μm^3 to 1 mm^3 is a volume increase of a thousandfold. That's why Cox is obsessively focused on automating every step of the process, from the rats' video training to the tracing of the connectome. "These IARPA projects force scientific research to look a lot more like engineering," he says. "We need to turn the crank fast."

Speeding up the experiments allows Cox's team to test more theories regarding the brain's circuits, which will help the AI researchers too. In machine learning, computer scientists set the overall architecture of the neural network while the program itself decides how to connect its many computations into sequences. So the researchers plan to train the rats and a neural network on the same visual-recognition task and compare both the patterns of links and the outcomes. "If we see certain connectivity motifs in the brain and we don't see them in the models, maybe that's a hint that we're doing something wrong," Cox says.

One area of investigation involves the brain's learning rules. Object recognition is thought to occur through a hierarchy of processing, with a

THE FUTURE according to...

Ray Kurzweil

Cofounder and chancellor,
Singularity University

When will we have computers as capable as the brain?
I believe computers will match and then quickly exceed human capabilities in the areas where humans are still superior today by 2029.

How will brainlike computers change the world?
AI is already accelerating our ability to find cures for diseases, improve crop yields and other forms of productivity to reduce poverty, and find solutions to environmental problems. It already represents a brain extender. Who today can do their job or get an education without these extensions to our intelligence?

Do you have any qualms about a future in which computers have human-level (or greater) intelligence?
Every technology since fire has had intertwined promise and peril. I believe that our best strategy to keep AI safe and beneficial is to essentially merge with it. We are already on that path. Whether [intelligent] machines are inside or outside [the] body is not a critical issue. But they will ultimately go inside our bodies and brains (because miniaturization is another exponential trend), keeping us healthy (by augmenting our immune system), providing virtual and augmented reality from within the nervous system, and making us smarter.

PHOTO-ILLUSTRATION BY Gluekit

first set of neurons taking in the basics of color and form, the next set finding edges to separate the object from its background, and so on. As the animal gets better at a recognition task, researchers can ask: Which set of neurons in the hierarchy is changing its activity most drastically? And as the AI gets better at the same task, does the pattern of activity in its neural network change in the same way as the rat's?

IARPA hopes the findings will apply not only to computer vision but also to machine learning in general. "There's a bit of a leap of faith that we all take here, but I think it's an evidence-based leap of faith," Cox says. He notes that the cerebral cortex, the outer layer of neural tissue where high-level cognition occurs, has a "suspiciously similar" structure throughout. To neuroscientists and AI experts, that uniformity suggests that a fundamental type of circuit may be used throughout the brain for information processing, which they hope to discover. Defining that ur-circuit may be a step toward a generally intelligent AI.

●■●

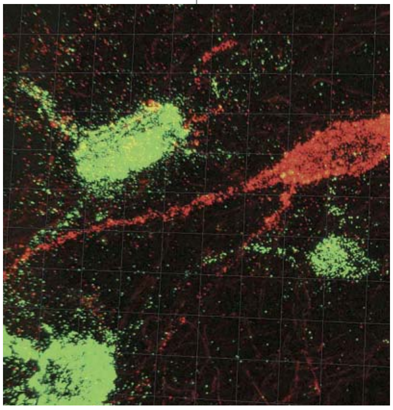

FLUORESCENT IN SITU SEQUENCING

Each neuron gets an "RNA bar code" consisting of a unique sequence of molecules called bases. The sequencing machine reads the bar codes by recording the differently colored flashes of light associated with the four types of bases. George Church's team maps where these bar codes appear in the brain tissue to show where the neurons connect.

TONY ZADOR/COLD SPRING HARBOR LABORATORY

WHILE COX'S TEAM is turning the crank, attempting to make tried-and-true neuroscience procedures run faster, another Microns researcher is pursuing a radical idea. If it works, says George Church, a professor at the Wyss Institute for Biologically Inspired Engineering, at Harvard University, it could revolutionize brain science.

Church is coleading a Microns team with Tai Sing Lee of Carnegie Mellon University, in Pittsburgh. Church is responsible for the connectome-mapping part of the process, and he's taking a drastically different approach from the other teams. He doesn't use electron microscopy to trace axon connections; he believes the technique is too slow and produces too many errors. As the other teams try to trace axons across a cubic millimeter of tissue, Church says, errors will accumulate and muddy the connectome data.

Church's method isn't affected by the length of axons or the size of the brain chunk under investigation. He uses genetically engineered mice and a technique called DNA bar coding, which tags each neuron with a unique genetic identifier that can be read out from the fringy tips of its dendrites to the terminus of its long axon. "It doesn't matter if you have some gargantuan long axon," he says. "With bar coding you find the two ends, and it doesn't matter how much confusion there is along the way." His team uses slices of brain tissue that are thicker than those used by Cox's team–20 μm instead of 30 nm–because they don't have to worry about losing the path of an axon from one slice to the next. DNA sequencing machines record all the bar codes present in a given slice of brain tissue, and then a program sorts through the genetic information to make a map showing which neurons connect to one another.

Church and his collaborator Anthony Zador, a neuroscience professor at Cold Spring Harbor Laboratory, in New York, have proven in prior experiments that the bar coding and sequencing technique works, but they haven't yet put the data together into the connectome maps that the Microns project requires. Assuming his team gets it done, Church says Microns will mark just the beginning of his brain-mapping efforts: He next wants to chart all the connections of an entire mouse brain, with its 70 million neurons and 70 billion connections. "Doing 1 cubic millimeter is extraordinarily myopic," Church says. "My ambition doesn't end there."

Such large-scale maps could provide insights for developing AIs that more rigorously mimic biological brains. But Church, who relishes the role of provocateur, envisions another way forward for computing: Stop trying to build silicon copies of brains, he says, and instead build biological brains that are even better at handling the computational tasks that human brains are so good at. "I think we'll soon have the ability to do synthetic neurobiology, to actually build brains that are variations on natural brains," he says. While silicon-based computers beat biological systems when it comes to processing speeds, Church imagines engineered brains augmented with circuit elements to accelerate their operations.

In Church's estimation, the Microns goal of reverse engineering the brain may not be achievable. The brain is so complex, he says, that even if researchers succeed in building these machines, they may not fully understand the brain's mysteries–and that's okay. "I think understanding is a bit of a fetish among scientists," Church says. "It may be much easier to engineer the brain than to understand it." ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/iarpamicrons0617

A Road Map for the Artificial Brain

LARGE-SCALE BRAINLIKE SYSTEMS ARE POSSIBLE WITH EXISTING TECHNOLOGY—IF WE'RE WILLING TO SPEND THE MONEY

By **Jennifer Hasler**

BRAIN-INSPIRED COMPUTING is having a moment. Artificial neural network algorithms like deep learning, which are very loosely based on the way the human brain operates, now allow digital computers to perform such extraordinary feats as translating language, hunting for subtle patterns in huge amounts of data, and beating the best human players at Go. ● But even as engineers continue to push this mighty computing strategy, the energy efficiency of digital computing is fast approaching its limits. Our data centers and supercomputers already draw megawatts—some 2 percent of the electricity consumed in the United States goes to data centers alone. The human brain, by contrast, runs quite well on about 20 watts, which represents the power produced by just a fraction of the food a person eats each day. If we want to keep improving computing, we will need our computers to become more like our brains.

PHOTOGRAPH BY Dan Saelinger

Hence the recent focus on neuromorphic technology, which promises to move computing beyond simple neural networks and toward circuits that operate more like the brain's neurons and synapses do. The development of such physical brainlike circuitry is actually pretty far along. Work at my lab and others around the world over the past 35 years has led to artificial neural components like synapses and dendrites that respond to and produce electrical signals much like the real thing.

So, what would it take to integrate these building blocks into a brain-scale computer? In 2013, Bo Marr, a former graduate student of mine at Georgia Tech, and I looked at the best engineering and neuroscience knowledge of the time and concluded that it should be possible to build a silicon version of the human cerebral cortex with the transistor technology then in production. What's more, the resulting machine would take up less than a cubic meter of space and consume less than 100 watts, not too far from the human brain.

That is not to say creating such a computer would be easy. The system we envisioned would still require a few billion dollars to design and build, including some significant packaging innovations to make it compact. There is also the question of how we would program and train the computer. Neuromorphic researchers are still struggling to understand how to make thousands of artificial neurons work together and how to translate brainlike activity into useful engineering applications.

Still, the fact that we can envision such a system means that we may not be far off from smaller-scale chips that could be used in portable and wearable electronics. These gadgets demand low power consumption, and so a highly energy-efficient neuromorphic chip–even if it takes on only a subset of computational tasks, such as signal processing–could be revolutionary. Existing capabilities, like speech recognition, could be extended to handle noisy environments. We could even imagine future smartphones conducting real-time language translation between you and the person you're talking to. Think of it this way: In the 40 years since the first signal-processing integrated circuits, Moore's Law has improved energy efficiency by roughly a factor of 1,000. The most brainlike neuromorphic chips could dwarf such improvements, potentially driving down power consumption by another factor of 100 million. That would bring computations that would otherwise need a data center to the palm of your hand.

* * *

THE ULTIMATE BRAINLIKE MACHINE will be one in which we build analogues for all the essential functional components of the brain: the synapses, which connect neurons and allow them to receive and respond to signals; the dendrites, which combine and perform local computations on those incoming signals; and the core, or soma, region of each neuron, which integrates inputs from the dendrites and transmits its output on the axon.

Simple versions of all these basic components have already been implemented in silicon. The starting point for such work is the same metal-oxide-semiconductor field-effect transistor, or MOSFET, that is used by the billions to build the logic circuitry in modern digital processors.

These devices have a lot in common with neurons. Neurons operate using voltage-controlled barriers, and their electrical and chemical activity depends primarily on channels in which ions move between the interior and exterior of the cell–a smooth, analog process that involves a steady buildup or decline instead of a simple on-off operation.

MOSFETs are also voltage controlled and operate by the movement of individual units of charge. And when MOSFETs are operated in the "subthreshold" mode, below the voltage threshold used to digitally switch between on and off, the amount of current flowing through the device is very small–less than a thousandth of what is

THE FUTURE according to...

Carver Mead

Professor emeritus, California Institute of Technology

Do you have any qualms about a future in which computers have human-level (or greater) intelligence? Every time there's been progress in technology, people have predicted that it will be the end of society as we know it, and it never has been. The world today is a vastly better world for everyone than the world of 100 years ago. In fact, technology has been the single force that has propagated prosperity across the world. Technology has always been used [for both good and evil], and yet if you look over 100 years, the good has won out over the bad by a large margin.

PHOTO-ILLUSTRATION BY Gluekit

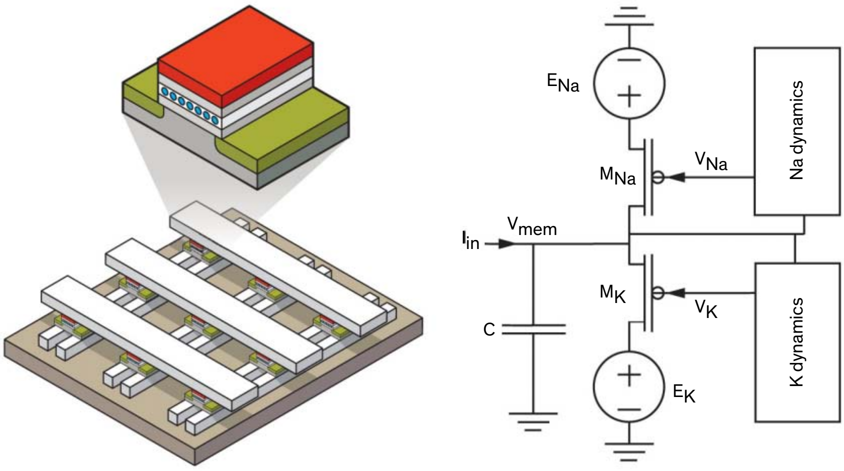

SYNAPSES AND SOMA: The floating-gate transistor [top left], which can store differing amounts of charge, can be used to build a "crossbar" array of artificial synapses [bottom left]. Electronic versions of other neuron components, such as the soma region [right], can be made from standard transistors and other circuit components.

seen in the typical switching of digital logic gates.

The notion that subthreshold transistor physics could be used to build brainlike circuitry originated with Carver Mead of Caltech, who helped revolutionize the field of very-large-scale circuit design in the 1970s. Mead pointed out that chip designers fail to take advantage of a lot of interesting behavior—and thus information—when they use transistors only for digital logic. The process, he wrote in 1990, essentially involves "taking all the beautiful physics that is built into...transistors, mashing it down to a 1 or 0, and then painfully building it back up with AND and OR gates to reinvent the multiply." A more "physical" or "physics-based" computer could execute more computations per unit energy than its digital counterpart. Mead predicted such a computer would take up significantly less space as well.

In the intervening years, neuromorphic engineers have made all the basic building blocks of the brain out of silicon with a great deal of biological fidelity. The neuron's dendrite, axon, and soma components can all be fabricated from standard transistors and other circuit elements. In 2005, for example, Ethan Farquhar, then a Ph.D. candidate, and I created a neuron circuit using a set of six MOSFETs and a handful of capacitors. Our model generated electrical pulses that very closely matched those in the soma part of a squid neuron, a long-standing experimental subject. What's more, our circuit accomplished this feat with similar current levels and energy consumption to those in the squid's brain. If we had instead used analog circuits to model the equations neuroscientists have developed to describe that behavior, we'd need on the order of 10 times as many transistors. Performing those calculations with a digital computer would require even more space.

* * *

EMULATING SYNAPSES is a little trickier. A device that behaves like a synapse must have the ability to remember what state it is in, respond in a particular way to an incoming signal, and adapt its response over time.

There are a number of potential approaches to building synapses. The most mature one by far is the single-transistor learning synapse (STLS), a device that my colleagues and I at Caltech worked on in the 1990s while I was a graduate student studying under Mead.

We first presented the STLS in 1994, and it became an important tool for engineers who were building modern analog circuitry, such as physical neural networks. In neural networks, each node in the network has a weight associated with it, and those weights determine how data from different nodes are combined. The STLS was the first device that could hold a variety of different weights and be reprogrammed on the fly. The device is also nonvolatile, which means that it remembers its state even when not in use—a capability that significantly reduces how much energy it needs.

The STLS is a type of floating-gate transistor, a device that is used to build memory cells in flash memory. In an ordinary MOSFET, a gate controls the flow of electricity through a current-carrying channel. A floating-gate transistor has a second gate that sits between this electrical gate and the channel. This floating gate is not directly connected to ground or any other component. Thanks to that electrical isolation, which is enhanced by high-quality silicon-insulator interfaces, charges remain in the floating gate for a long time. The floating gate can take on many different amounts of charge and so have many different levels of electrical response, an essential requisite for creating an artificial synapse capable of varying its response to stimuli.

My colleagues and I used the STLS to demonstrate the first crossbar network, a computational model currently popular with nanodevice researchers. In this two-dimensional array, devices sit at the intersection of input lines running north-south and output lines running east-west. This configuration is useful because it lets you program the con-

ILLUSTRATION BY James Provost

nection strength of each "synapse" individually, without disturbing the other elements in the array.

Thanks in part to a recent Defense Advanced Research Projects Agency program called SyNAPSE, the neuromorphic engineering field has seen a surge of research into artificial synapses built from nanodevices such as memristors, resistive RAM, and phase-change memories (as well as floating-gate devices). But it will be hard for these new artificial synapses to improve on our two-decade-old floating-gate arrays. Memristors and other novel memories come with programming challenges; some have device architectures that make it difficult to target a single specific device in a crossbar array. Others need a dedicated transistor in order to be programmed, adding significantly to their footprint. Because floating-gate memory is programmable over a wide range of values, it can be more easily fine-tuned to compensate for manufacturing variation from device to device than can many nanodevices. A number of neuromorphic research groups that tried integrating nanodevices into their designs have recently come around to using floating-gate devices.

* * *

SO HOW WILL WE PUT all these brainlike components together? In the human brain, of course, neurons and synapses are intermingled. Neuromorphic chip designers must take a more integrated approach as well, with all neural components on the same chip, tightly mixed together. This is not the case in many neuromorphic labs today: To make research projects more manageable, different building blocks may be placed in different areas. Synapses, for example, may be relegated to an off-chip array. Connections may be routed through another chip called a field-programmable gate array, or FPGA.

But as we scale up neuromorphic systems, we'll need to take care that we don't replicate the arrangement in today's computers, which lose a significant amount of energy driving bits back and forth between logic, memory, and storage. Today, a computer can easily consume 10 times the energy to move the data needed for a multiple-accumulate operation—a common signal-processing computation—as it does to perform the calculation.

The brain, by contrast, minimizes the energy cost of communication by keeping operations highly local. The memory elements of the brain, such as synaptic strengths, are mixed in with the neural components that integrate signals. And the brain's "wires"—the dendrites and axons that extend from neurons to transmit, respectively, incoming signals and outgoing pulses—are generally fairly short relative to the size of the brain, so they don't require large amounts of energy to sustain a signal. From anatomical data, we know that more than 90 percent of neurons connect with only their nearest 1,000 or so neighbors.

Another big question for the builders of brainlike chips and computers is the algorithms we will run on them. Even a loosely brain-inspired system can have a big advantage over digital systems. In 2004, for example, my group used floating-gate devices to perform multiplications for signal processing with just 1/1,000 the energy and 1/100 the area of a comparable digital system. In the years since, other researchers and my group have successfully demonstrated neuromorphic approaches to many other kinds of signal-processing calculations.

But the brain is still 100,000 times as efficient as the systems in these demonstrations. That's because while our current neuromorphic technology takes advantage of the neuronlike physics of transistors, it doesn't consider the algorithms the brain uses to perform its operations.

Today, we are just beginning to discover these physical algorithms—that is, the processes that will allow brainlike chips to operate with more brainlike efficiency. Four years ago, my research group used silicon somas, synapses, and dendrites to perform a word-spotting algorithm that identifies words in a speech waveform. This physical algorithm exhibited a thousandfold improvement in energy efficiency over predicted analog signal processing. Eventually, by lowering the amount of voltage supplied to the chips and using smaller transistors, researchers should be able to build chips that parallel the brain in efficiency for a range of computations.

When I started in neuromorphic research 30 years ago, everyone believed tremendous opportunities would arise from designing systems that are more like the brain. And indeed, entire industries are now being built around brain-inspired AI and deep learning, with applications that promise to transform—among other things—our mobile devices, our financial institutions, and how we interact in public spaces.

And yet, these applications rely only slightly on what we know about how the brain actually works. The next 30 years will undoubtedly see the incorporation of more such knowledge. We already have much of the basic hardware we need to accomplish this neuroscience-to-computing translation. But we must develop a better understanding of how that hardware should behave—and what computational schemes will yield the greatest real-world benefits.

Consider this a call to action. We have come pretty far with a very loose model of how the brain works. But neuroscience could lead to far more sophisticated brainlike computers. And what greater feat could there be than using our own brains to learn how to build new ones? ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/brainroadmap0617

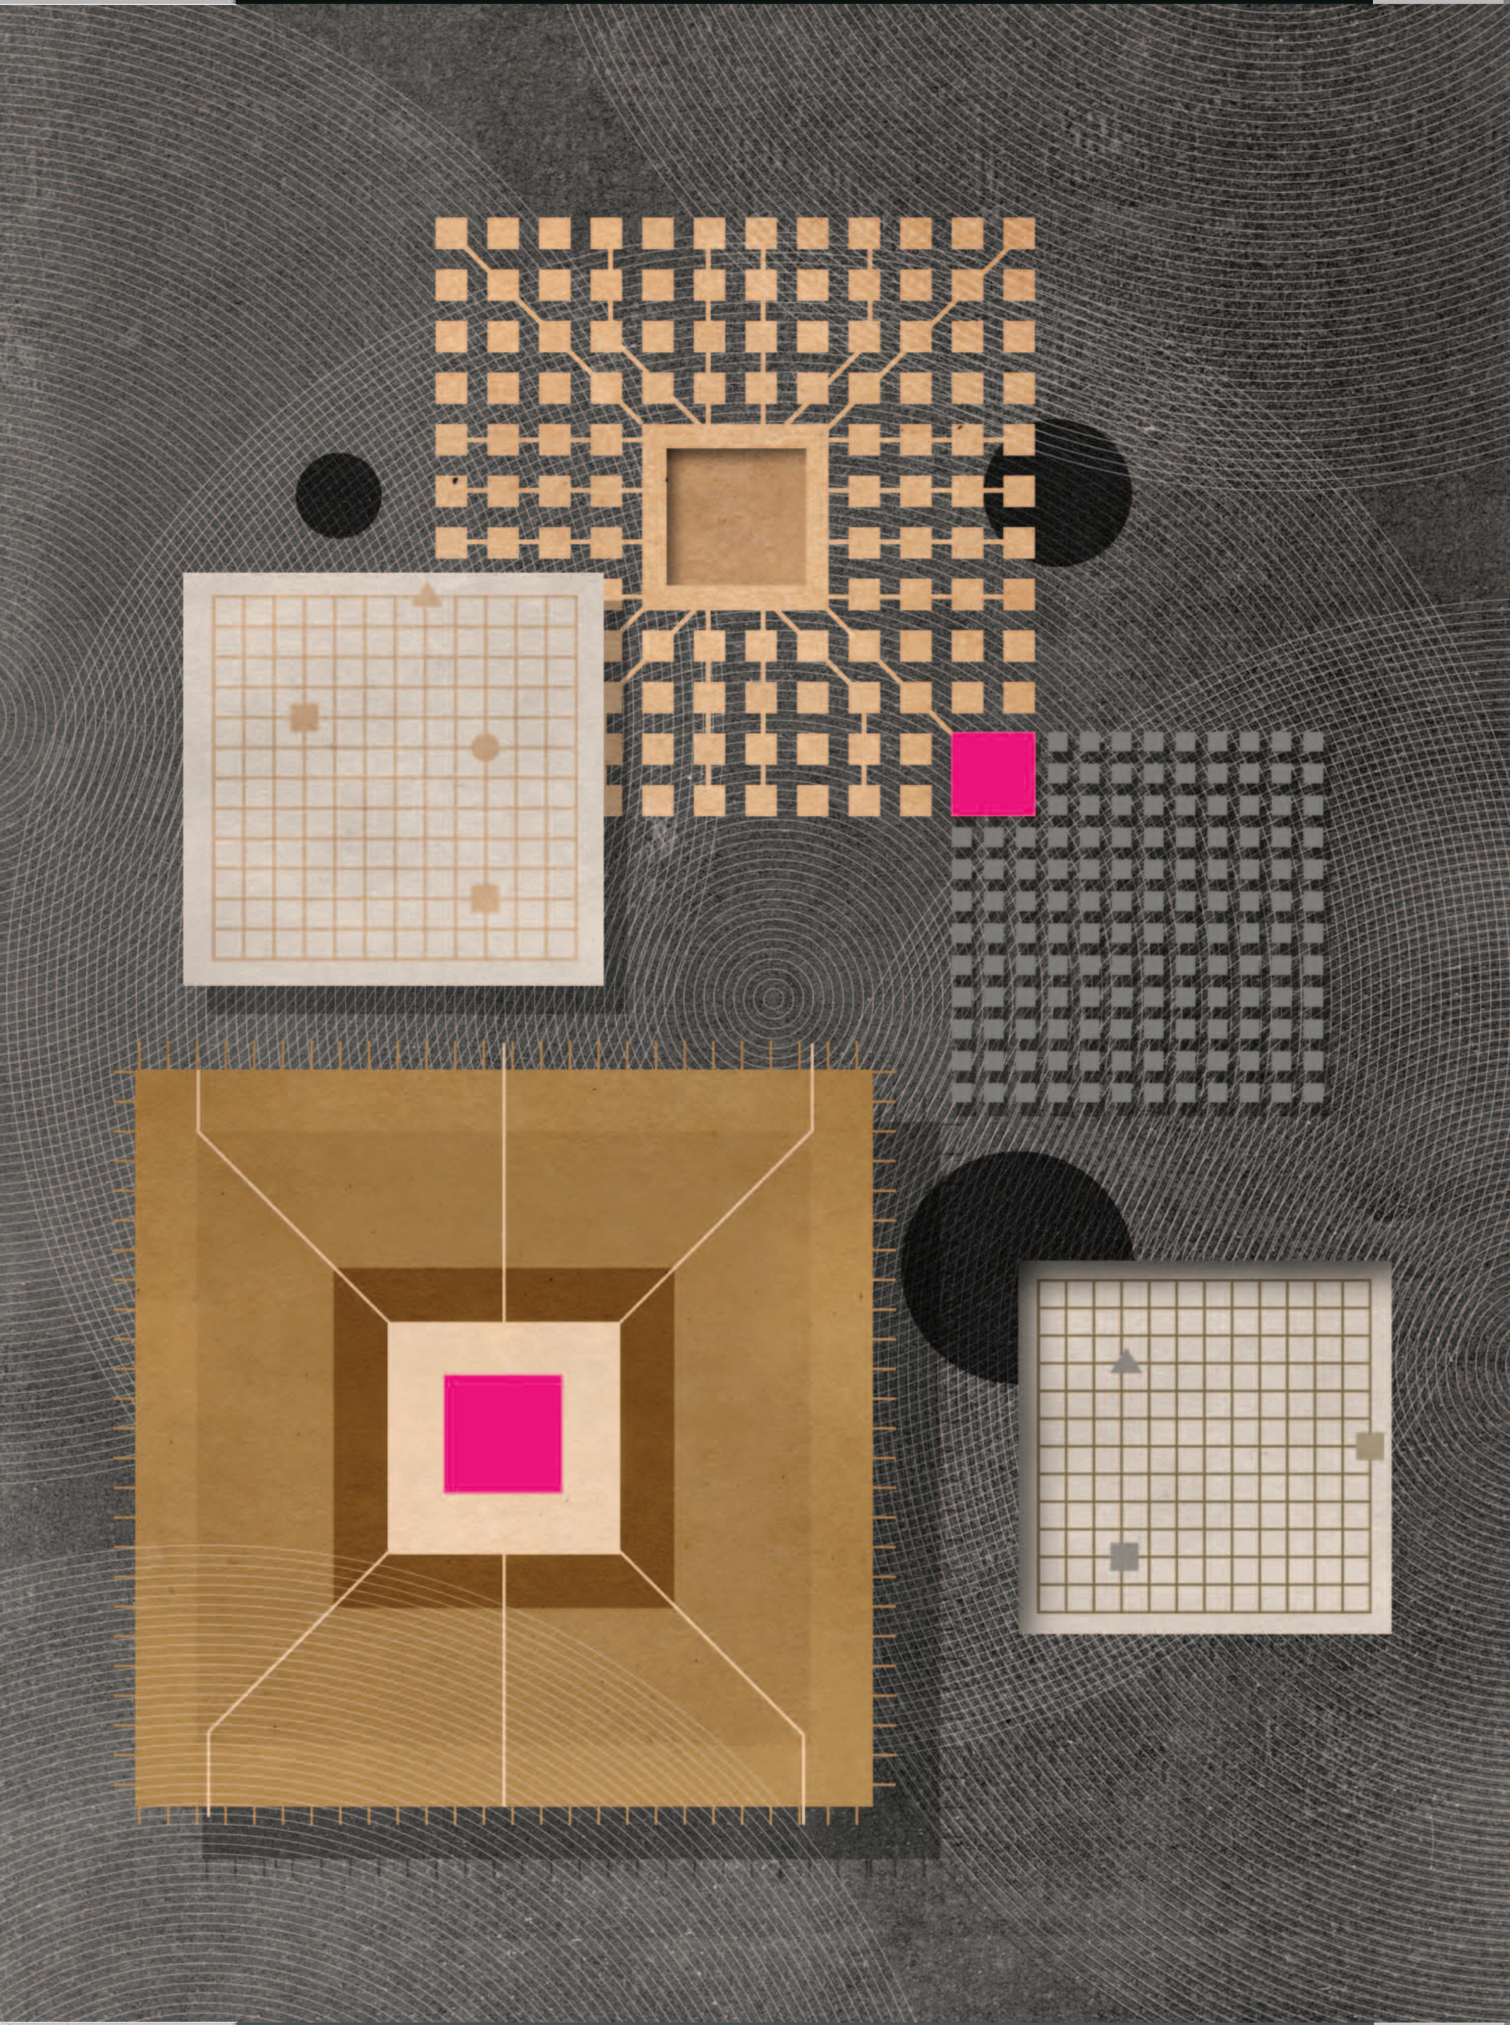

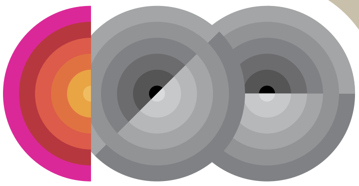

The Neuromorphic Chip's Make-or-Break Moment

COULD DEEP LEARNING BE THIS TECHNOLOGY'S KILLER APP?

By **Lee Gomes**

PEOPLE IN THE TECH WORLD talk of a technology "crossing the chasm" by making the leap from early adopters to the mass market. A case study in chasm crossing is now unfolding in neuromorphic computing. ● The approach mimics the way neurons are connected and communicate in the human brain, and enthusiasts say neuromorphic chips can run on much less power than traditional CPUs. The problem, though, is proving that neuromorphics can move from research labs to commercial applications. The field's leading researchers spoke frankly about that challenge at the Neuro Inspired Computational Elements Workshop, held in March at the IBM research facility at Almaden, Calif. ● "There currently is a lot of hype about neuromorphic computing," said Steve Furber, the researcher at the University of Manchester, in England, who heads the SpiNNaker project, a major neuromorphics effort. "It's true that neuromorphic systems exist, and

ILLUSTRATION BY Chad Hagen

you can get one and use one. But all of them have fairly small user bases, in universities or industrial research groups. All require fairly specialized knowledge. And there is currently no compelling demonstration of a high-volume application where neuromorphic outperforms the alternative."

Other attendees gave their own candid analyses. Another prominent researcher, Chris Eliasmith of the University of Waterloo, in Ontario, Canada, said the field needs to meet the hype issue "head-on." Given that neuromorphics has generated a great deal of excitement, Eliasmith doesn't want to "fritter it away on toy problems": A typical neuromorphic demonstration these days will show a system running a relatively simple artificial intelligence application. Rudimentary robots with neuromorphic chips have navigated down a Colorado mountain trail and rolled over squares of a specific color placed in a pattern on the floor. The real test is for traditional companies to accept neuromorphics as a mainstay platform for everyday engineering challenges, Eliasmith said, but there is "tons more to do" before that happens.

* * *

THE BASIC BUILDING BLOCK of neuromorphic computing is what researchers call a spiking neuron, which plays a role analogous to what a logic gate does in traditional computing. In the central processing unit of your desktop, transistors are assembled into different types of logic gates—AND, OR, XOR, and the like—each of which evaluates two binary inputs. Then, based on those values and the gate's type, each gate outputs either a 1 or a 0 to the next logic gate in line. All of them work in precise synchronization to the drumbeat of the chip's master clock, mirroring the Boolean logic of the software it's running.

The spiking neuron is a different beast. Imagine a node sitting on a circuit and measuring whatever spikes—in the form of electrical pulses—are transmitted along the circuit. If a certain number of spikes occur within a certain period of time, the node is programmed to send along one or more new spikes of its own, the exact number depending on the design of the particular chip. Unlike the binary, 0-or-1 option of traditional CPUs, the responses to spikes can be weighted to a range of values, giving neuromorphics something of an analog flavor. The chips save on energy in large part because their neurons aren't constantly firing, as occurs with traditional silicon technology, but instead become activated only when they receive a spiking signal.

A neuromorphic system connects these spiking neurons into complex networks, often according to a task-specific layout that programmers have worked out in advance. In a network designed for image recognition, for example, certain connections between neurons take on certain weights, and the way spikes travel between these neurons with their respective weights can be made to represent different objects. If one pattern of spikes appears at the output, programmers would know the image is of a cat; another pattern of spikes would indicate the image is of a chair.

Within neuromorphics, each research group has come up with its own design to make this possible. IBM's DARPA-funded TrueNorth neuromorphic chip, for example, does its spiking in custom hardware, while Furber's SpiNNaker (Spiking Neural Network Architecture) relies on software running on the ARM processors that he helped develop.

THE FUTURE according to...

Nick Bostrom

Author of Superintelligence: Paths, Dangers, Strategies

When will we have computers as capable as the brain?
Hardware-wise, the brain still compares favorably with machines. Estimates vary, but perhaps the cortex performs something like 10^{16} or 10^{18} operations per second using 20 watts, which is impressive. Eventually, the limits of computation in machine substrate are of course far beyond those in biological tissue, and it shouldn't take too long to reach rough equivalence. The advance of algorithms is harder to predict, but the notion that we could have human-level AI within a small number of decades seems credible, though there is great uncertainty on both the lower and upper sides of this estimate.

What are your biggest qualms about a future in which computers have human-level (or greater) intelligence?
I think it's a whole new era that dawns at that point, and it is hard to foretell in much detail what it might contain. But one dimension of concern is that we fail to solve the problem of scalable control: how to engineer advanced artificial intelligence such that it will continue to behave as intended and act as an extension of human will, even as its intelligence is increased to arbitrarily high levels. This is an active area of research at the moment.

PHOTO-ILLUSTRATION BY Gluekit

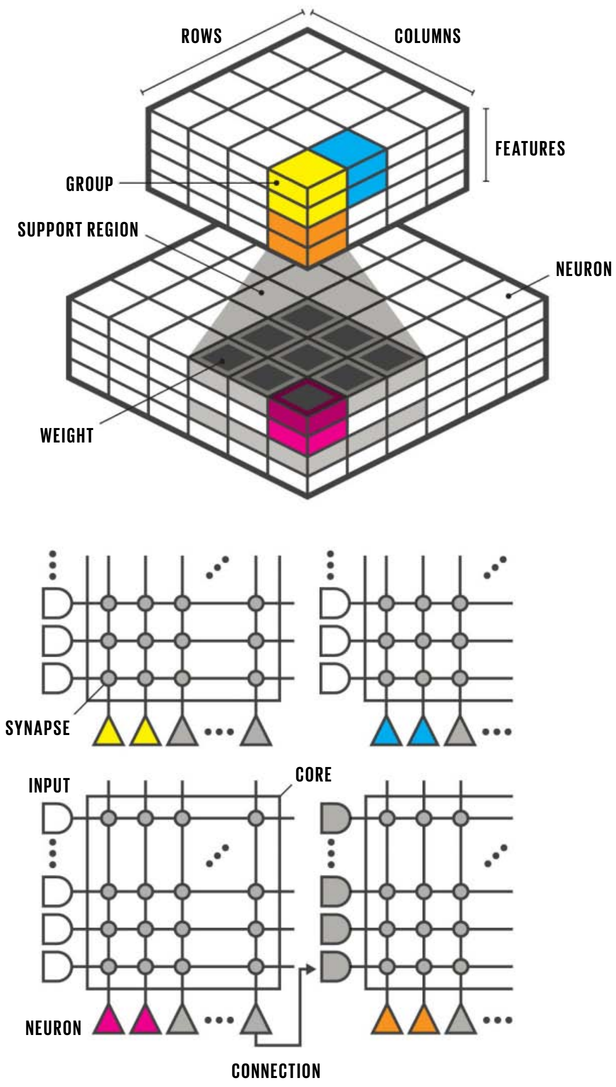

TINY SPIKES

Two layers within a neural network contain groups of "neurons" with similar functions, indicated by color [blue, yellow, orange, and pink] in the top illustration. In the bottom graphic, those neurons are mapped to spiking neurons in an IBM TrueNorth chip. The spiking neurons are connected by gridlike "synapses" to other neurons in the same core, and to a row of inputs. Those inputs can generate spikes, which are then processed by the neural network.

* * *

IN THE EARLY DAYS, there was no consensus on what neuromorphic systems would actually do, except to somehow be useful in brain research. In truth, spiking chips were something of a solution looking for a problem. Help, though, arrived unexpectedly from an entirely different part of the computing world.

ILLUSTRATION BY James Provost

ENGINEERING COGNITION 3

Starting in the 1990s, artificial intelligence researchers made a number of theoretical advances involving the design of the "neural networks" that had been used for decades for computational problem solving, though with limited success. Emre Neftci, with the University of California, Irvine's Neuromorphic Machine Intelligence Lab, said that when combined with faster silicon chips, these new, improved neural networks allowed computers to make dramatic advances in classic computing problems, such as image recognition.

This new breed of computing tools used what's come to be called deep learning, and in the past few years, deep learning has basically taken over the computer industry. Members of the neuromorphics research community soon discovered that they could take a deep-learning network and run it on their new style of hardware. And they could take advantage of the technology's power efficiency: The TrueNorth chip, which is the size of a postage stamp and holds a million "neurons," is designed to use a tiny fraction of the power of a standard processor.

Those power savings, say neuromorphics boosters, will take deep learning to places it couldn't previously go, such as inside a mobile phone, and into the world's hottest technology market. Today, deep learning enables many of the most widely used mobile features, such as the speech recognition required when you ask Siri a question. But the actual processing occurs on giant servers in the cloud, for lack of sufficient computing horsepower on the device. With neuromorphics on board, say its supporters, everything could be computed locally.

Which means that neuromorphic computing has, to a considerable degree, hitched its wagon to deep learning's star. When IBM wanted to show off a killer app for its TrueNorth chip, it ran a deep neural network that classified images. Much of the neuromorphics community now defines success as being able to supply extremely power-efficient chips for deep learning, first for big server farms such as those run by Google, and later for mobile phones and other small, power-sensitive applications. The former is considered the easier engineering challenge, and neuromorphics optimists say commercial products for server farms could show up in as few as two years.

* * *

UNFORTUNATELY FOR NEUROMORPHICS, just about everyone else in the semiconductor industry—including big players like Intel and Nvidia—also wants in on the deep-learning market. And that market might turn out to be one of the rare cases in which the incumbents, rather than the innovators, have the strategic advantage. That's because deep learning, arguably the most advanced software on the planet, generally runs on extremely simple hardware.

Karl Freund, an analyst with Moor Insights & Strategy who specializes in deep learning, said the key bit of computation involved in running a deep-learning system–known as matrix multiplication–can easily be handled with 16-bit and even 8-bit CPU components, as opposed to the 32- and 64-bit circuits of an advanced desktop processor. In fact, most deep-learning systems use traditional silicon, especially the graphics coprocessors found in the video cards best known for powering video games. Graphics coprocessors can have thousands of cores, all working in tandem, and the more cores there are, the more efficient the deep-learning network.

So chip companies are bringing out deep-learning chips that are made out of very simple, traditional components, optimized to use as little power as possible. (That's true of Google's Tensor Processing Unit, the chip the search company announced last year in connection with its own deep-learning efforts.) Put differently, neuromorphics' main competition as the platform of choice for deep learning is an advanced generation of what are essentially "vanilla" silicon chips.

Some companies on the vanilla side of this argument deny that neuromorphic systems have an edge in power efficiency. William J. Dally, a Stanford electrical engineering professor and chief scientist at Nvidia, said that the demonstrations performed with TrueNorth used a very early version of deep learning, one with much less accuracy than is possible with more recent systems. When accuracy is taken into account, he said, any energy advantage of neuromorphics disappears.

"People who do conventional neural networks get results and win the competitions," Dally said. "The neuromorphic approaches are interesting scientifically, but they are nowhere close on accuracy."

Indeed, researchers have yet to figure out simple ways to get neuromorphic systems to run the huge variety of deep-learning networks that have been developed on conventional chips. Brian Van Essen, at the Center for Applied Scientific Computing at the Lawrence Livermore National Laboratory, said his group has been able to get neural networks to run on TrueNorth but that the task of picking the right network and then successfully porting it over remains "a challenge." Other researchers say the most advanced deep-learning systems require more neurons, with more possible interconnections, than current neuromorphic technology can offer.

● ■ ●

THE NEUROMORPHICS COMMUNITY must tackle these problems with a small pool of talent. The March conference, the field's flagship event, attracted only a few hundred people; meetings associated with deep learning usually draw many thousands. IBM, which declined to

THE FUTURE according to...

Rodney Brooks

Chairman and CTO, Rethink Robotics

When will we have computers as capable as the brain?
Rodney Brooks's revised question: When will we have computers/robots recognizably as intelligent and as conscious as humans? Not in our lifetimes, not even in Ray Kurzweil's lifetime, and despite his fervent wishes, just like the rest of us, he will die within just a few decades. It will be well over 100 years before we see this level in our machines. Maybe many hundred years.

As intelligent and as conscious as dogs? Maybe in 50 to 100 years. But they won't have noses anywhere near as good as the real thing. They will be olfactorily challenged dogs.

How will brainlike computers change the world?
Since we won't have intelligent computers like humans for well over 100 years, we cannot make any sensible projections about how they will change the world, as we don't understand what the world will be like at all in 100 years. (For example, imagine reading Turing's paper on computable numbers in 1936 and trying to project out how computers would change the world in just 70 or 80 years.) So an equivalent well-grounded question would have to be something simpler, like "How will computers/robots continue to change the world?" Answer: Within 20 years most baby boomers are going to have robotic devices in their homes, helping them maintain their independence as they age in place. This will include Ray Kurzweil, who will still not be immortal.

Do you have any qualms about a future in which computers have human-level (or greater) intelligence?
No qualms at all, as the world will have evolved so much in the next 100+ years that we cannot possibly imagine what it will be like, so there is no point in qualming. Qualming in the face of zero facts or understanding is a fun parlor game but generally not useful. And yes, this includes Nick Bostrom.

PHOTO-ILLUSTRATION BY Gluekit

comment for this article, said last fall that TrueNorth, which debuted in 2014, is now running experiments and applications for more than 130 users at more than 40 universities and research centers.

By contrast, there is hardly a Web company or university computer department on the planet that isn't doing something with deep learning on conventional chips. As a result, those conventional architectures have a robust suite of development tools, along with legions of engineers trained in their use—typical advantages of an incumbent technology with a large installed base. Getting the deep-learning community to switch to a new and unfamiliar way of doing things will prove extremely difficult unless neuromorphics can offer an unmistakable performance and power advantage.

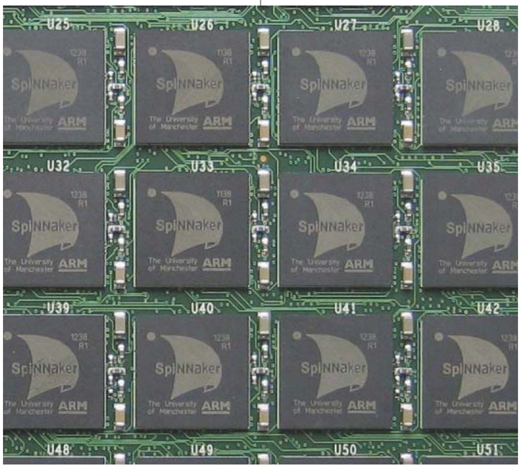

BUILDING BLOCKS: The SpiNNaker project is constructing a machine with 50,000 of these specialized chips in hopes of creating a network of 1 billion "neurons."

UNIVERSITY OF MANCHESTER

Again, that's a problem the neuromorphics community openly acknowledges. If the presentations at the March conference frequently referred to the challenges that lie ahead for the field, most of them also offered suggestions on how to overcome them.

The University of Waterloo's Eliasmith, for example, said that neuromorphics must progress on a number of fronts. One of them is building more-robust hardware, with more neurons and interconnections, to handle more-advanced deep-learning systems. Also needed, he said, are theoretical insights about the inherent strengths and weaknesses of neuromorphic systems, to better know how to use them most productively. To be sure, he still believes the technology can live up to expectations. "We have been seeing regular improvements, so I'm encouraged," Eliasmith said.

Still, the neuromorphics community might find that its current symbiotic relationship with deep learning comes with its own hazards. For all the recent successes of deep learning, plenty of experts still question how much of an advance it will turn out to be.

Deep learning clearly delivers superior results in applications such as pattern recognition, in which one picture is matched to another picture, or for language translation. It remains to be seen how far the technique will take researchers toward the holy grail of "generalized intelligence," or the ability of a computer to have, like HAL 9000 in the film *2001: A Space Odyssey*, the reasoning and language skills of a human. Deep-learning pioneer Yann LeCun compares AI research to driving in the fog. He says there is a chance that even armed with deep learning, AI might any day now crash into another brick wall.

That prospect caused some at the conference to suggest that neuromorphics researchers should persevere even if the technology doesn't deliver a home run for deep learning. Bruno Olshausen, director of the University of California, Berkeley's Redwood Center for Theoretical Neuroscience, said neuromorphic technology may, on its own, someday bring about AI results more sophisticated than anything deep learning ever could. "When we look at how neurons compute in the brain, there are concrete things we can learn," he said. "Let's try to build chips that do the same thing, and see what we can leverage out of them."

The SpiNNaker project's Furber echoed those sentiments when asked to predict when neuromorphics would be able to produce low-power components that could be used in mobile phones. His estimate was five years—but he said he was only 80 percent confident in that prediction. He added, however, that he was far more certain that neuromorphics would play an important role in studying the brain, just as early proponents thought it might.

However, there is a meta-issue hovering over the neuromorphics community: Researchers don't know whether the spiking behavior they are mimicking in the brain is central to the way the mind works, or merely one of its many accidental by-products. Indeed, the surest way to start an argument with a neuromorphics researcher is to suggest that we don't really know enough about how the brain works to have any business trying to copy it in silicon. The usual response you'll get is that while we certainly don't know everything, we clearly know enough to start.

It has often been noted that progress in aviation was made only after inventors stopped trying to copy the flapping wings of birds and instead discovered—and then harnessed—basic forces, such as thrust and lift. The knock against neuromorphic computing is that it's stuck at the level of mimicking flapping wings, an accusation the neuromorphics side obviously rejects. Depending on who is right, the field will either take flight and soar over the chasm, or drop into obscurity. ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/neuromorphic0617

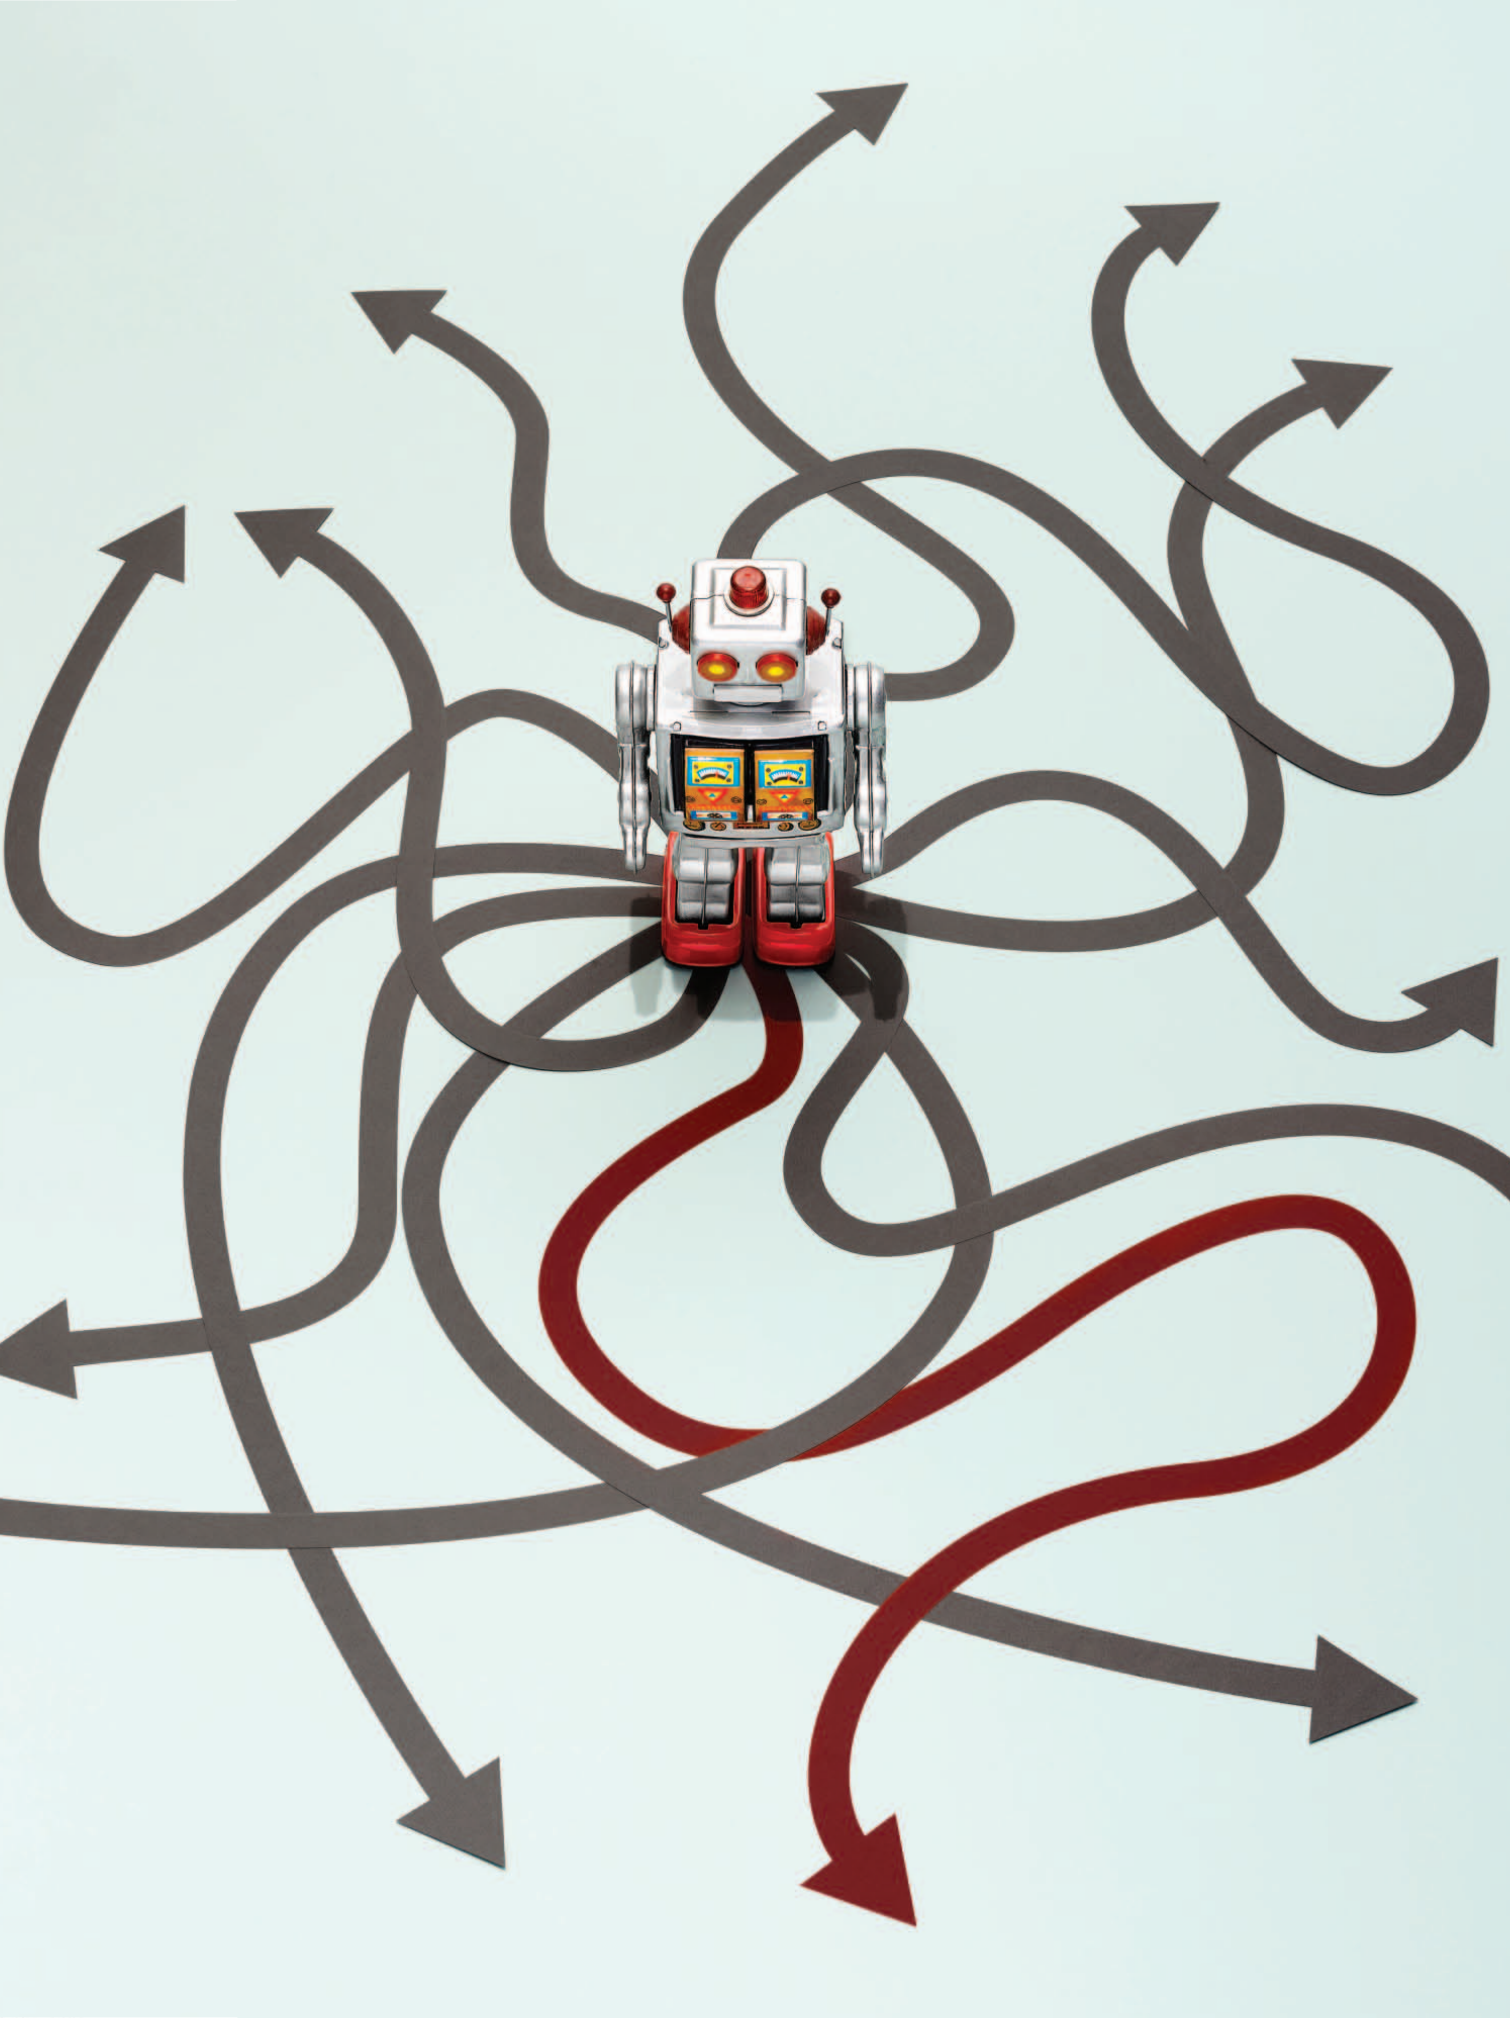

Navigate Like a Rat

RODENTS ARE WORLD-CLASS PATHFINDERS, SO WHY NOT BUILD ROBOTS WITH RODENT-LIKE BRAINS?

By Jean Kumagai

IF YOU TAKE A COMMON BROWN RAT and drop it into a lab maze or a subway tunnel, it will immediately begin to explore its surroundings, sniffing around the edges, brushing its whiskers against surfaces, peering around corners and obstacles. After a while, it will return to where it started, and from then on, it will treat the explored terrain as familiar. ● Roboticists have long dreamed of giving their creations similar navigation skills. To be useful in our environments, robots must be able to find their way around on their own. Some are already learning to do that in homes, offices, warehouses, hospitals, hotels, and, in the case of self-driving cars, entire cities. Despite the progress, though, these robotic platforms still struggle to operate reliably under even mildly challenging conditions. Self-driving vehicles, for example, may come equipped with sophisticated sensors and detailed maps of the road ahead, and yet human drivers still

PHOTOGRAPH BY Dan Saelinger

have to take control in heavy rain or snow, or at night.

The lowly brown rat, by contrast, is a nimble navigator that has no problem finding its way around, under, over, and through the toughest spaces. When a rat explores an unfamiliar territory, specialized neurons in its 2-gram brain fire, or spike, in response to landmarks or boundaries. Other neurons spike at regular distances—once every 20 centimeters, every meter, and so on—creating a kind of mental representation of space. Yet other neurons act like an internal compass, recording the direction in which the animal's head is turned. Taken together, this neural activity allows the rat to remember where it's been and how it got there. Whenever it follows the same path, the spikes strengthen, making the rat's navigation more robust.

So why can't a robot be more like a rat?

The answer is, it can. At the Queensland University of Technology (QUT), in Brisbane, Australia, Michael Milford and his collaborators have spent the last 14 years honing a robot navigation system modeled on the brains of rats. This biologically inspired approach, they hope, could help robots navigate dynamic environments without requiring advanced, costly sensors and computationally intensive algorithms.

An earlier version of their system allowed an indoor package-delivery bot to operate autonomously for two weeks in a lab. During that period, it made more than 1,100 mock deliveries, traveled a total of 40 kilometers, and recharged itself 23 times. Another version successfully mapped an entire suburb of Brisbane, using only the imagery captured by the camera on a MacBook. Now Milford's group is translating its rat-brain algorithms into a rugged navigation system for the heavy-equipment maker Caterpillar, which plans to deploy it on a fleet of underground mining vehicles.

* * *

MILFORD, WHO'S 35 and looks about 10 years younger, began investigating brain-based navigation in 2003, when he was a Ph.D. student at the University of Queensland working with roboticist Gordon Wyeth, who's now dean of science and engineering at QUT.

At the time, one of the big pushes in robotics was the "kidnapped robot" problem: If you take a robot and move it somewhere else, can it figure out where it is? One way to solve the problem is SLAM, which stands for simultaneous localization and mapping. While running a SLAM algorithm, a robot can explore strange terrain, building a map of its surroundings while at the same time positioning, or localizing, itself within that map.

Wyeth had long been interested in brain-inspired computing, starting with work on neural networks in the late 1980s. And so he and Milford decided to work on a version of SLAM that took its cues from the rat's neural circuitry. They called it RatSLAM.

There already were numerous flavors of SLAM, and today they number in the dozens, each with its own advantages and drawbacks. What they all have in common is that they rely on two separate streams of data. One relates to what the environment looks like, and robots gather this kind of data using sensors as varied as sonars, cameras, and laser scanners. The second stream concerns the robot itself, or more specifically, its speed and orientation; robots derive that data from sensors like rotary encoders on their wheels or an inertial measurement unit (IMU) on their bodies. A SLAM algorithm looks at the environmental data and tries to identify notable landmarks, adding these to its map. As the robot moves, it monitors its speed and direction and looks for those landmarks; if the robot recognizes a landmark, it uses the landmark's position to refine its own location on the map.

But whereas most implementations of SLAM aim for highly detailed, static maps, Milford and Wyeth were more interested in how to navigate through an environment that's in constant flux. Their aim wasn't to create maps built with costly lidars and high-powered computers—they wanted their system to make sense of space the way animals do.

"Rats don't build maps," Wyeth says. "They have other ways of remembering where they are." Those ways include neurons called place cells and head-direction cells, which respectively let the rat identify landmarks and gauge its direction. Like other neurons, these cells are densely interconnected and work by adjusting their spiking patterns in response to different stimuli. To mimic this structure and behavior in software, Milford adopted a type of artificial neural network called an attractor network. These neural nets consist of hundreds to thousands of interconnected nodes that, like groups of neurons, respond to an input by producing a specific spiking pattern, known as an attractor state. Computational neuroscientists use attractor networks to study neurons associated with memory and motor behavior. Milford and Wyeth wanted to use them to power RatSLAM.

They spent months working on the software, and then they loaded it into a Pioneer robot, a mobile platform popular among roboticists. Their rat-brained bot was alive.

But it was a failure. When they let it run in a 2-by-2-meter arena, Milford says, "it got lost even in that simple environment."

* * *

MILFORD AND WYETH REALIZED that RatSLAM didn't have enough information with which to reduce errors as it made its decisions. Like other SLAM algorithms, it doesn't try

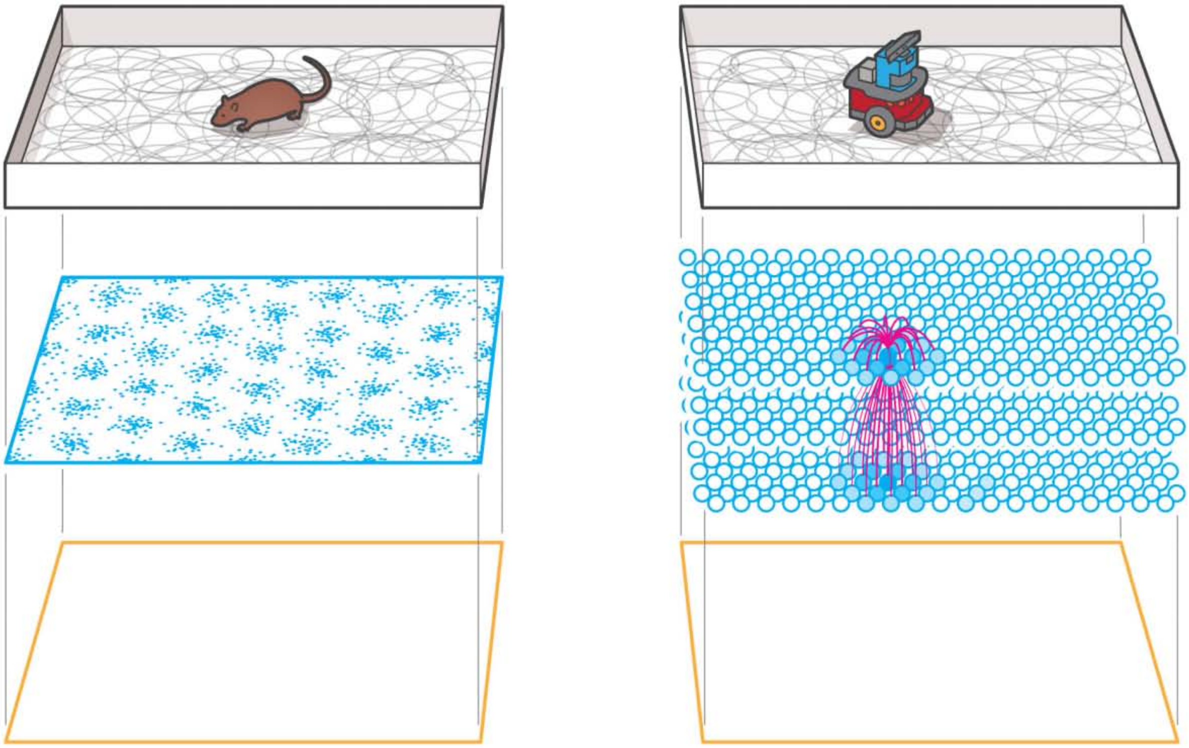

RAT RACE

As a rat explores an unfamiliar arena [top level, left], the spiking of neurons called grid cells [shown as blue dots, middle left] reveals a regular hexagonal pattern [bottom left]. The RatSLAM algorithm mimics a rodent's navigation neurons. As a robot running RatSLAM moves around a new space [top level, right], interconnected nodes in the algorithm's attractor neural network respond with a spiking pattern [middle right] much as grid cells do, forming a similarly regular pattern [bottom right].

to make exact, definite calculations about where things are on the map it's generating; instead, it relies on approximations and probabilities as a way of incorporating uncertainties—conflicting sensor readings, for example—that inevitably crop up. If you don't take that into account, your robot ends up lost.

That seemed to be the problem with RatSLAM. In some cases, the robot would recognize a landmark and be able to refine its position, but other times the data was too ambiguous. After not too long, the accrued error was bigger than 2 meters—the robot thought it was outside the arena!

In other words, their rat-brain model was too crude. It needed better neural circuitry to be able to abstract more information about the world.

"So we engineered a new type of neuron, which we called a 'pose' cell," Milford says. The pose cell didn't just tell the robot its location or its orientation, it did both at the same time. Now, when the robot identified a landmark it had seen before, it could more precisely encode its place on the map and keep errors in check.

Again, Milford placed the robot inside the 2-by-2-meter arena. "Suddenly, our robot could navigate quite well," he recalls.

Interestingly, not long after the researchers devised these artificial cells, neuroscientists in Norway and the United Kingdom announced the discovery of grid cells, which are neurons whose spiking activity forms regular geometric patterns and tells the animal its relative position within a certain area.

"Our pose cells weren't exactly grid cells, but they had similar features," Milford says. "That was rather gratifying."

The robot tests moved to bigger arenas with greater complexity. "We did a whole floor, then multiple floors in the building," Wyeth recalls. "Then I told Michael, 'Let's do a whole suburb.' I thought he would kill me."

Milford loaded the RatSLAM software into a MacBook and taped it on the roof of his red 1994 Mazda Astina. To get a stream of data about the environment, he used the laptop's camera, setting it to snap a photo of the street ahead of the car several times per second. To get a stream of the data about the robot itself—in this case, his car—he found a creative solution. Instead of attaching encoders to the wheels

ILLUSTRATION BY James Provost

or using an IMU or GPS, he used simple image-processing techniques. By tracking and comparing pixels on sequences of photos from the MacBook, his SLAM algorithm could calculate the vehicle's speed as well as direction changes.

Milford drove for about 2 hours through the streets of the Brisbane suburb of St. Lucia, covering 66 kilometers. The result wasn't a precise, to-scale map, but it accurately represented the topology of the roads and could pinpoint exactly where the car was at any given moment. RatSLAM worked.

"It immediately drew attention and was widely discussed because it was very different from what other roboticists were doing," says David Wettergreen, a roboticist at Carnegie Mellon University, in Pittsburgh, who specializes in autonomous robots for planetary exploration. Indeed, it's still considered one of the most notable examples of brain-inspired robotics.

But though RatSLAM created a stir, it didn't set off a wave of research based on those same principles. And when Milford and Wyeth approached companies about commercializing their system, they found many keen to hear their pitch but ultimately no takers. "A colleague told me we should have called it 'NeuroSLAM,' " Wyeth says. "People have bad associations with rats."

That's why Milford is excited about the two-year project with Caterpillar, which began in March. "I've always wanted to create systems that had real-world uses," he says. "It took a lot longer than I expected for that to happen."

● ■ ●

"WE LOOKED AT THEIR RESULTS and decided this is something we could get up and running quickly," Dave Smith, an engineer at Caterpillar's Australia Research Center, in Brisbane, tells me. "The fact that it's rat inspired is just a cool thing."

Underground mines are among the harshest man-made places on earth. They're cold, dark, and dusty, and due to the possibility of a sudden collapse or explosion, they're also extremely dangerous. For companies operating in such an extreme environment, improving their ability to track machines and people underground is critical.

In a surface mine, you'd simply use high-precision differential GPS, but that obviously doesn't work below ground. Existing indoor navigation systems, such as laser mapping and RF networks, are expensive and often require infrastructure that's difficult to deploy and maintain in the severe conditions of a mine. For instance, when Caterpillar engineers considered 3D lidar, like the ones used on self-driving cars, they concluded that "none of them can survive underground," Smith says.

One big reason that mine operators need to track their vehicles is to plan how they excavate. Each day starts with a "dig plan" that specifies the amount of ore that will be mined in various tunnels. At the end of the day, the operator compares the dig plan to what was actually mined, to come up with the next day's dig plan. "If you're feeding in inaccurate information, your plan is not going to be very good. You may start mining dirt instead of ore, or the whole tunnel could cave in," Smith explains. "It's really important to know what you've done."

The traditional method is for the miner to jot down his movements throughout the day, but that means he has to stop what he's doing to fill out paperwork, and he's often guessing what actually occurred. The QUT navigation system will more accurately measure where and how far each vehicle travels, as well as provide a reading of where the vehicle is at any given time. The first vehicle

THE FUTURE according to...

Gary Marcus

Professor of psychology,
New York University

PHOTO-ILLUSTRATION BY **Gluekit**

When will we have computers as capable as the brain?
Computers are already far more capable than brains in many respects (for example, arithmetic and memory), but I think it could still be 20 to 50 years before machines have the ability to read and comprehend and reason about novel situations as fluently as people can.

How will brainlike computers change the world?
If machines can read and reason at human levels while being able to compute at superhuman levels, we may see unprecedented progress in science and medicine.

Do you have any qualms about a future in which computers have human-level (or greater) intelligence?
Sure, but my biggest worry is about machines having too much power, not about them being too smart. You could, for example, be president of the United States and do a lot of damage, regardless of what your IQ is.

How far will machine learning take us?
Current techniques aren't strong enough to solve hard problems in language and common sense, but eventually we will have more powerful techniques, perhaps fairly different from what we have now, and the sky will be the limit when we do.

will drive into the mine and map the environment using the rat-brain-inspired navigation algorithm, while also gathering images of each tunnel with a low-cost 720p camera. The only unusual feature of the camera is its extreme ruggedization, which Smith says goes well beyond military specifications.

Subsequent vehicles will use those results to localize themselves within the mine, comparing footage from their own cameras with previously gathered images. The vehicles won't be autonomous, Milford notes, but that capability could eventually be achieved by combining the camera data with data from IMUs and other sensors. This would add more precision to the trucks' positioning, allowing them to drive themselves.

The QUT team has started collecting data within actual mines, which will be merged with another large data set from Caterpillar containing about a thousand hours of underground camera imagery. They will then devise a preliminary algorithm, to be tested in an abandoned mine somewhere in Queensland, with the help of Mining3, an Australian mining R&D company; the Queensland government is also a partner on the project. The system could be useful for deep open-pit mines, where GPS tends not to work reliably. If all goes well, Caterpillar plans to commercialize the system quickly. "We need these solutions," Smith says.

● ■ ●

FOR NOW, MILFORD'S TEAM relies on standard computing hardware to run its algorithms, although they keep tabs on the latest research in neuromorphic computing. "It's still a bit early for us to dive in," Milford says. Eventually, though, he expects his brain-inspired systems will map well to neuromorphic chip architectures like IBM's True North and the University of Manchester's SpiNNaker. [For more on these chips, see "The Neuromorphic Chip's Make-or-Break Moment," in this issue.]

Will brain-inspired navigation ever go mainstream? Many developers of self-driving cars, for instance, invest heavily in creating detailed maps of the roads where their vehicles will drive. The vehicles then use their cameras, lidars, GPS, and other sensors to locate themselves on the maps, rather than having to build their own.

Still, autonomous vehicles need to prove they can drive in conditions like heavy rain, snow, fog, and darkness. They also need to better handle uncertainty in the data; images with glare, for instance, might have contributed to a fatal accident involving a self-driving Tesla last year. Some companies are already testing machine-learning-based navigation systems, which rely on artificial neural networks, but it's possible that more brain-inspired approaches like RatSLAM could complement those systems, improving performance in difficult or unexpected scenarios.

Carnegie Mellon's Wettergreen offers a more tantalizing possibility: giving cars the ability to navigate to specific locations without having to explicitly plan a trajectory on a city map. Future robots, he notes, will have everything modeled down to the millimeter. "But I don't," he says, "and yet I can still find my way around. The human brain uses different types of models and maps–some are metric, some are more topological, and some are semantic."

A human, he continues, can start with an idea like "Somewhere on the south side of the city, there's a good Mexican restaurant." Arriving in that general area, the person can then look for clues as to where the restaurant may be. "Even the most capable self-driving car wouldn't know what to do with that kind of task, but a more brain-inspired system just might."

Some roboticists, however, are skeptical that such unconventional approaches to SLAM are going to pay off. As sensors like lidar, IMUs, and GPS get better and cheaper, traditional SLAM algorithms will be able to produce increasingly accurate results by combining data from multiple sources. People tend to ignore the fact that "SLAM is really a sensor fusion problem and that we are getting better and better at doing SLAM with lower-cost sensors," says Melonee Wise, CEO of Fetch Robotics, a company based in San Jose, Calif., that sells mobile robots for transporting goods in highly dynamic environments. "I think this disregard causes people to fixate on trying to solve SLAM with one sensor, like a camera, but in today's low-cost sensor world that's not really necessary."

Even if RatSLAM doesn't become practical for most applications, developing such brainlike algorithms offers us a window into our own intelligence, says Peter Stratton, a computer scientist at the Queensland Brain Institute who collaborates with Milford. He notes that standard computing's von Neumann architecture, where the processor is separated from memory and data is shuttled between them, is very inefficient.

"The brain doesn't work anything like that. Memory and processing are both happening in the neuron. It's 'computing with memories,' " Stratton says. A better understanding of brain activity, not only as it relates to responses to stimuli but also in terms of its deeper internal processes–memory retrieval, problem solving, daydreaming–is "what's been missing from past AI attempts," he says.

Milford notes that a lot of types of intelligence aren't easy to study using only animals. But when you observe how rats and robots perform the same tasks, like navigating a new environment, you can test your theories about how the brain works. You can replay scenarios repeatedly. You can tinker and manipulate your models and algorithms. "And unlike with an animal or an insect brain," he says, "we can see everything in a robot's 'brain.' " ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/ratslam0617

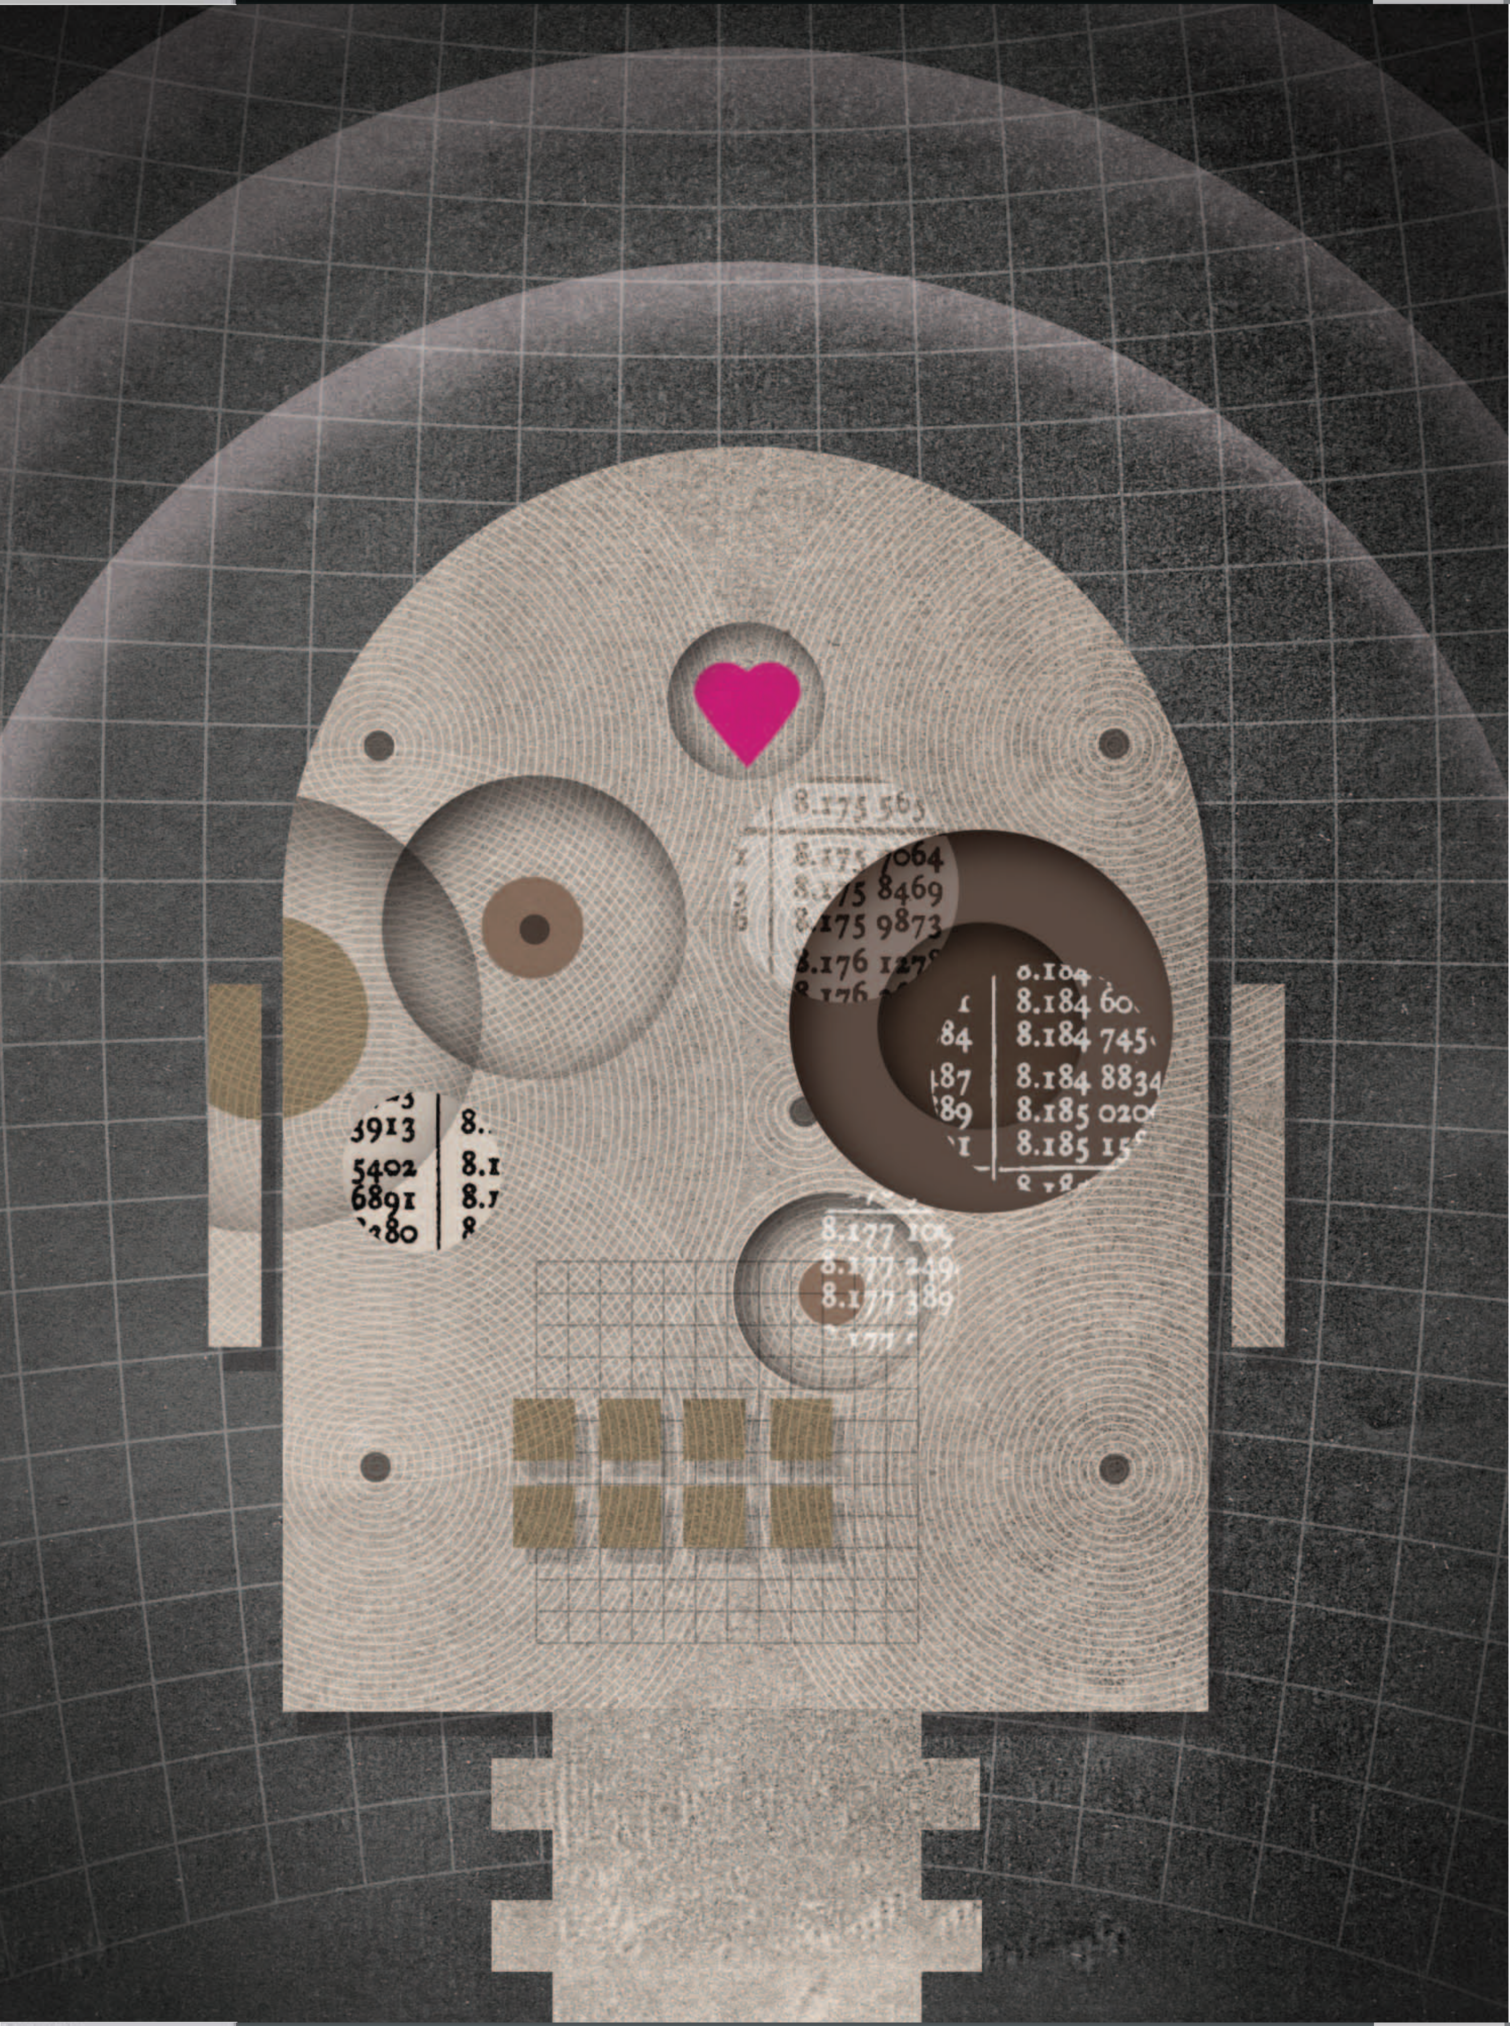
8.175 565
8.175 7064
8.175 8469
8.175 9873
8.176 127
8.184 60
8.184 745
8.184 8834
8.185 020
8.185 15
3913
5402
6891
8.177 105
8.177 249
8.177 389

Can We Quantify Machine Consciousness?

BRAINLIKE CIRCUITRY MIGHT ONE DAY ENDOW SOME COMPUTERS WITH AWARENESS. HERE'S HOW WE'D KNOW

By Christof Koch & Giulio Tononi

IMAGINE THAT AT SOME time in the not-too-distant future, you've bought a smartphone that comes bundled with a personal digital assistant (PDA) living in the cloud. You assign a sexy female voice to the PDA and give it access to all of your emails, social media accounts, calendar, photo album, contacts, and other bits and flotsam of your digital life. She—for that's how you quickly think of her—knows you better than your mother, your soon-to-be ex-wife, your friends, or your therapist. Her command of English is flawless; you have endless conversations about daily events; she gets your jokes. She is the last voice you hear before you drift off to sleep and the first upon awakening. You panic when she's off-line. She becomes indispensable to your well-being and so, naturally, you fall in love. Occasionally, you wonder whether she truly reciprocates your feelings and whether she is even capable of experiencing

ILLUSTRATION BY Chad Hagen

anything at all. But the warm, husky tone of her voice and her ability to be that perfect foil to your narcissistic desires overcome these existential doubts. Alas, your infatuation eventually cools off after you realize she is carrying on equally intimate conversations with thousands of other customers.

This, of course, is the plot of *Her*, a 2013 movie in which an anodyne Theodore Twombly falls in love with the software PDA Samantha.

Over the next few decades such a fictional scenario will become real and commonplace. Deep machine learning, speech recognition, and related technologies have dramatically progressed, leading to Amazon's Alexa, Apple's Siri, Google's Now, and Microsoft's Cortana. These virtual assistants will continue to improve until they become hard to distinguish from real people, except that they'll be endowed with perfect recall, poise, and patience—unlike any living being.

The availability of such digital simulacra of many qualities we consider uniquely human will raise profound scientific, psychological, philosophical, and ethical questions. These emulations will ultimately upend the way we think about ourselves, about human exceptionalism, and about our place in the great scheme of things.

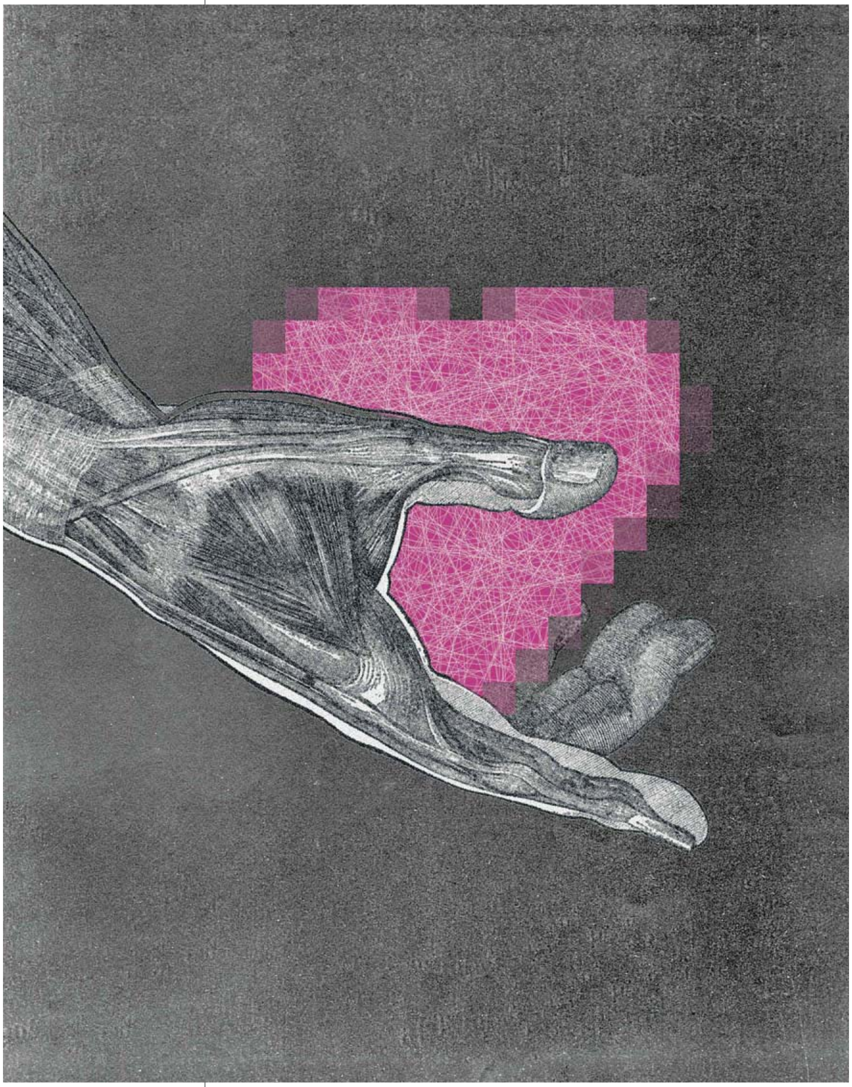

Here we will survey the intellectual lay of the land concerning these coming developments. Our view is that as long as such machines are based on present-day computer architectures, they may act just like people—and we may be tempted to treat them that way—but they will, in fact, feel nothing at all. If computers are built more like the brain is, though, they could well achieve true consciousness.

* * *

THE FAITH OF OUR AGE is faith in the digital computer—programmed properly, it will give us all we wish. Cornucopia. Indeed, smart money in Silicon Valley holds that digital computers will be able to replicate and soon exceed anything and everything that humans are capable of.

But could sufficiently advanced computers ever become conscious? One answer comes from those who subscribe to computationalism, the reigning theory of mind in contemporary philosophy, psychology, and neuroscience. It avers that all mental states—such as your conscious experience of a god-awful toothache or the love you feel for your partner—are computational states. These are fully characterized by their functional relationships to relevant sensory inputs, behavioral outputs, and other computational states in between. That is, brains are elaborate input-output devices that compute and process symbolic representations of the world. Brains are computers, with our minds being the software.

ILLUSTRATION BY **Chad Hagen**

Adherents to computationalism apply these precepts not only to brains and to the behavior they generate but also to the way it *feels* to be a brain in a particular state. After all, that's what consciousness is: any subjective feeling, any experience—what we see, hear, feel, remember, think.

Computationalism assumes that my painful experience of a toothache is but a state of my brain in which certain nerve cells are active in response to the infected tooth, leading to my propensity to moan, hold my jaw, not eat on that side of my mouth, inability to focus on other tasks, and so on. If all of these states are simulated in software on a digital computer, the thinking goes, the system as a whole will not only behave exactly like me but also *feel and think* exactly like me. That is, consciousness is computable. Explicitly or implicitly, this is one of the central tenets held by the digerati in academe, media, and industry.

In this view, there is nothing more to consciousness than the instantiation of the relevant computational states. Nothing else matters, including how the computations are implemented physically, whether on the hardware of a digital computer or on the squishy stuff inside the skull. According to computationalism, a future Samantha—or even better, an embodied example like Ava in the brilliant, dark movie *Ex Machina*—will have experiences and feelings just as we do. She will experience sights and sounds, pleasure and pain, love and hate.

Or perhaps she won't.

Computationalism is based on the assumption that if two systems are *functionally* indistinguishable, they will be *mentally* indistinguishable. Because we experience the world, the argument goes, a digital computer that is functionally equivalent to us would necessarily also experience the world as we do—that is, it would also be conscious. But is this assumption warranted? To answer such a question, we need a principled, quantitative theory of what consciousness is and what it takes for a physical system to have it.

● ■ ●

UNTIL RECENTLY, such a theory of consciousness wasn't available. True, neuroscientists like us have been engaged in the difficult search for the "neural correlates of consciousness," carrying out increasingly elaborate experiments on people and related species such as monkeys and mice. These experiments have identified regions in the neocortex, the outer surface of the brain just underneath the skull, that are critically involved in consciously seeing and hearing things. Yet having been directly involved in this empirical research program, we know that even if such a quest proves reasonably successful, identifying some particular brain structures or modes of neural activity necessary for consciousness in people or closely related animals will not be sufficient to establish whether creatures with very different nervous systems—such as an octopus or a bee—*are* conscious, by how much, or of what. And any such discovery in neuroscience will be insufficient to establish whether or not machines can be conscious.

There is, however, a fundamental theory of consciousness that offers hope for a principled answer to the question of consciousness in entities vastly different from us, including machines. That theory does not start from behavior or from the brain. Instead, it begins from consciousness itself—from our own experience, the only one we are absolutely certain of. This is the bedrock of certainty that René Descartes, father of modern philosophy, science, and analytic geometry, referred to in the most famous deduction in Western thought: *I think, therefore I am*.

This theory, called integrated information theory, or IIT, has been developed over the past two decades. It attempts to define what consciousness is, what it takes for a physical system to have it, and how one can measure, at least in principle, both its quantity and its quality, starting from its physical substrate.

IIT is too involved for us to explain here; we can only sketch its general outlines. The theory identifies five essential properties that are true of every conceivable experience of consciousness: (1) Every experience exists intrinsically (for the subject of that experience, not for an external observer); (2) each experience is structured (it is composed of parts and the relations among them); (3) it is integrated (it cannot be subdivided into independent components); (4) it is definite (it has borders, including some contents and excluding others); and (5) it is specific (every experience is the way it is, and thereby different from trillions of possible others).

IIT then translates these properties into requirements that must be satisfied by any physical substrate for it to support consciousness. These requirements can be expressed mathematically and employed to assess the quantity and quality of consciousness for any physical system, whether it is the brain of a human, an octopus, or a bee—or the circuit board of a digital computer.

Crucially, according to IIT, the overall degree of consciousness does not depend on what the system does. Rather, it depends on how it is built—how it's physically put together. And only certain kinds of physical systems have the right kind of internal architecture to support consciousness: those that have a maximum of intrinsic cause-effect power, the causal power to determine their own states. In essence, this means that the system must be composed of many parts, each having specific causal powers within the overall system (the "information" part of IIT), and yet the system as a whole must not be reducible to those parts (the "integrated" part of IIT), making it far more powerful than the sum of its many parts.

IIT does not use the word "information" in its contemporary sense, as in "messages that are being passed by a sender to a receiver." Consciousness is not about information sent from one part of the brain to another. Instead, IIT refers to "information" in its original sense, with its root *inform*, meaning "to give form to." The power of any one mechanism, such as a brain or a computer, to influence its own next state, its causal power, gives rise to a form, a high-dimensional structure, that *is* experience.

IIT can explain in a principled manner many puzzling features of the neuroanatomy of consciousness—for instance, why the cerebellum, the little brain underneath the much bigger and better known neocortex, does not contribute to consciousness despite its having four times as many neurons: Its internal architecture, parallel sheets of feed-forward chains of neurons without much recurrent excitation, is very different from the highly heterogeneous, rich, and dense connectivity of the neocortex, which supports vast coalitions of active neurons that quickly assemble and disassemble. It also explains why consciousness fades during certain stages of sleep even though neocortical neurons continue to fire: Parts of the neocortex lose the ability to influence one another effectively.

IIT makes a number of counterintuitive predictions amenable to empirical tests. One prediction is that a nearly silent neocortex, in which few neurons are actively firing, has conscious experiences. Also, IIT has allowed Tononi and Marcello Massimini, now a professor at the University of Milan, to develop a device for assessing consciousness in humans, a combination of a magnetic coil to stimulate the brain and a high-density net of EEG electrodes to detect its response—a crude kind of consciousness meter. This device has already been used to ascertain whether brain-damaged or anesthetized patients unable to communicate are conscious or not.

Being a formal, mathematical theory, IIT can be applied to any physical system, be it the brain—a structure that evolved by natural selection—or an electronic circuit designed by engineers. As ongoing research shows, the physical architecture of certain parts of the neocortex—especially in the back, the way the neurons are connected—is ideal for maximizing the brain's intrinsic cause-effect power, its ability to be affected by its recent state and to determine its future state, which is why it supports consciousness.

By contrast, the physical architecture of a typical digital computer is absolutely inadequate, with very low connectivity at the gate level of its central processing unit and bottlenecks that prevent even a modicum of the necessary integration. That is, a computer may implement computations and functions judged to be intelligent from the perspective of a user looking at its output, but,

THE FUTURE according to...

Jürgen Schmidhuber

Scientific director, Swiss AI Lab IDSIA

When will we have computers as capable as the brain?
Soon. Every five years computing is getting roughly 10 times cheaper. Unlike Moore's Law, which says that the number of transistors per microchip doubles every 18 months (and which recently broke) this older trend has held since Konrad Zuse built the first working program-controlled computer. His machine could perform roughly one floating-point operation per second. Today, 75 years later, hardware is roughly a million billion times faster per unit price. Soon we'll have cheap devices with the raw computational power of a human brain; a few decades later, of all 10 billion human brains together, which collectively probably cannot execute more than 10^{30} meaningful elementary operations per second.

How will brainlike computers change the world?
Most current AI research is about using artificial neural networks to build friendly AIs that make their users healthier and happier and more addicted to their smartphones. But we already know how to implement principles of curiosity and creativity in self-motivated AIs pursuing their own goals. This will scale.

What will supersmart AIs do? Space is hostile to humans but friendly to appropriately designed robots, and offers many more resources than the thin film of biosphere around the Earth, which gets less than a billionth of the sun's light. While some AIs will remain fascinated with life, at least as long as they don't fully understand it, most will be more interested in the incredible new opportunities for robots and software life out there in space. Through innumerable self-replicating robot factories in the asteroid belt and elsewhere they will transform the rest of the solar system, and then within a few million years the entire galaxy, and within billions of years the rest of the reachable universe, held back only by the light speed limit.

PHOTO-ILLUSTRATION BY Gluekit

given its wiring, its intrinsic causal powers as a whole are minute compared with those of any brain. And this is true even if we treat the computer at a coarser level than transistors and resistors.

And here's the rub: That intrinsic power, the physical power to make a difference to oneself, cannot be computed or simulated. It has to be built into the physics of the system. A perfectly executed, biophysically accurate computer simulation of the human brain, including every one of its 86 billion neurons and its matrix of trillions of synapses, wouldn't be conscious. Even if this computer was hooked up to a speech synthesizer and told you about its supposed experiences, it would be nothing but behavior and functions cleverly executing programming. The beating heart of consciousness would be absent.

● ■ ●

THIS CONSEQUENCE OF IIT has sobering implications for those who hope that digital brain uploads may make people immortal. Their vision is that, within the next few decades, we will be able to accurately reconstruct the wiring scheme, the so-called connectome, of any one individual human brain and simulate it on appropriate digital hardware. This process would probably be destructive because there may be no way to access the brain's ultrastructure except by cutting it into wafer-thin slivers. Nevertheless, before you succumb to some deadly disease, you would upload a high-resolution version of your brain to the cloud. As long as the cloud infrastructure is up and running, your digital simulacrum will live on and on, interacting with other digital avatars. Rapture for nerds!

Yet, per IIT, this belief is as illusory as the belief in the hereafter of preceding prophets and religions. Although your digital simulacrum might speak and act as you would, it would be a complete zombie, experiencing nothing. Ironically, though, to your friends and loved ones back in the real world, you would have successfully transitioned to a sublime form of existence and would entice others to join you in an afterlife that, in fact, doesn't exist.

Whether or not IIT is correct is not merely of academic interest. Barring some global catastrophe, our society will create, within decades, machines with human-level intelligence and behaviors, able to understand speech and talk in many different languages, remember the past and anticipate the future, imagine novel scenarios, write books, compose music, direct films, conceive new goals, as well as move, drive, fly, and, inevitably, fight. From there, thanks to the availability of big data, the power of deep learning, and the speed of computing, it will be a short step to overcoming human limits. The birth of true artificial intelligence will profoundly affect mankind's future, including whether it has one.

Whether you are among those who believe that the arrival of human-level AI signals the dawn of paradise or the sunset of the age of humans, you will still have to answer a fundamental question: Are these AIs conscious? Does it *feel* like anything to be them? Or are they immensely more accomplished versions of present-day garbage disposal units, washing machines, or cars–extraordinarily clever machines, yes, but without sentience or feelings?

The answer to this question matters to the way we relate to future machines. If you take a hammer to your shiny Tesla, your friends might consider you crazy for destroying such a costly car; yet there is no question that you are free to do so. Try the same with your dog, though, and the police would rightfully arrest you. That's because a car is just a means to an end–a convenient way to get around town–while a dog is an end in itself, with some minimal rights, because it shares with us the gift of consciousness.

Finding the correct answer, however, cannot be left to our intuition. That might work temporarily for Theodore Twombly falling in love with Samantha, but given the gravity of the situation, we need guidance. We need a fundamental theory that specifies the exact conditions under which a particular system is capable of conscious experience.

IIT predicts that conventional digital computers running software will experience nothing like the movie we see and hear inside our heads. Because smart digital assistants and lifelike future robots are incapable of experience, as IIT insists, their software can be safely copied, edited, sold, pirated, or deleted. And they can be turned off, modified, destroyed, and replaced at will.

But the same need not be true for unconventional architectures. Special-purpose machines built following some of the same design principles as the brain, containing what's called neuromorphic hardware [see "The Neuromorphic Chip's Make-or-Break Moment," in this issue], could in principle be capable of substantial conscious experience. The key is that the logic and memory gates are heavily interconnected with a high degree of partially overlapping fan-in and fan-out between gates. (Compartmentalized components with highly specific functions do not contribute to intrinsic causal power.) The way the "brain" of the system is actually wired up, its (bio)physics, makes all the difference, not its input-output behavior.

Such a neuromorphic machine, if highly conscious, would then have intrinsic rights, in particular the right to its own life and well-being. In that case, society would have to learn to share the world with its own creations. ■

↗ POST YOUR COMMENTS at http://spectrum.ieee.org/consciousmachines0617

INTELLIGENT MACHINES

CONTINUED FROM PAGE 37

move three to four times a second, our visual perception is stable. He deduced that the brain must take account of how the eyes are moving; otherwise it would appear as if the world were wildly jumping about. Similarly, as you touch something, it would be confusing if the brain processed only the tactile input and didn't know how your fingers were moving at the same time. This principle of combining movement with changing sensations is called sensorimotor integration. How and where sensorimotor integration occurs in the brain is mostly a mystery.

Our discovery is that sensorimotor integration occurs in every region of the neocortex. It is not a separate step but an integral part of all sensory processing. Sensorimotor integration is a key part of the "intelligence algorithm" of the neocortex. We at Numenta have a theory and a model of exactly how neurons do this, one that maps well onto the complex anatomy seen in every neocortical region.

What are the implications of this discovery for machine intelligence? Consider two types of files you might find on a computer. One is an image file produced by a camera, and the other is a computer-aided design file produced by a program such as Autodesk. An image file represents a two-dimensional array of visual features. A CAD file also represents a set of features, but each feature is assigned a location in three-dimensional space. A CAD file models complete objects, not how the object appears from one perspective. With a CAD file, you can predict what an object will look like from any direction and determine how an object will interact with other 3D objects. You can't do these with an image file. Our discovery is that every region of the neocortex learns 3D models of objects much like a CAD program. Every time your body moves, the neocortex takes the current motor command, converts it into a location in the object's reference frame, and then combines the location with the sensory input to learn 3D models of the world.

In hindsight, this observation makes sense. Intelligent systems need to learn multidimensional models of the world. Sensorimotor integration doesn't occur in a few places in the brain; it is a core principle of brain function, part of the intelligence algorithm. Intelligent machines also must work this way.

● ■ ●

THESE THREE FUNDAMENTAL attributes of the neocortex–learning by rewiring, sparse distributed representations, and sensorimotor integration–will be cornerstones of machine intelligence. Future thinking machines can ignore many aspects of biology, but not these three. Undoubtedly, there will be other discoveries about neurobiology that reveal other aspects of cognition that will need to be incorporated into such machines in the future, but we can get started with what we know today.

From the earliest days of AI, critics dismissed the idea of trying to emulate human brains, often with the refrain that "airplanes don't flap their wings." In reality, Wilbur and Orville Wright studied birds in detail. To create lift, they studied bird-wing shapes and tested them in a wind tunnel. For propulsion, they went with a nonavian solution: propeller and motor. To control flight, they observed that birds twist their wings to bank and use their tails to maintain altitude during the turn. So that's what they did, too. Airplanes still use this method today, although we twist only the tail edge of the wings. In short, the Wright brothers studied birds and then chose which elements of bird flight were essential for human flight and which could be ignored. That's what we'll do to build thinking machines.

As I consider the future, I worry that we are not aiming high enough. While it is exciting for today's computers to classify images and recognize spoken queries, we are not close to building truly intelligent machines. I believe it is vitally important that we do so. The future success and even survival of humanity may depend on it. For example, if we are ever to inhabit other planets, we will need machines to act on our behalf, travel through space, build structures, mine resources, and independently solve complex problems in environments where humans cannot survive. Here on Earth, we face challenges related to disease, climate, and energy. Intelligent machines can help. For example, it should be possible to design intelligent machines that sense and act at the molecular scale. These machines would think about protein folding and gene expression in the same way you and I think about computers and staplers. They could think and act a million times as fast as a human. Such machines could cure diseases and keep our world habitable.

In the 1940s, the pioneers of the computing age sensed that computing was going to be big and beneficial, and that it would likely transform human society. But they could not predict exactly how computers would change our lives. Similarly, we can be confident that truly intelligent machines will transform our world for the better, even if today we can't predict exactly how. In 20 years, we will look back and see this as the time when advances in brain theory and machine learning started the era of true machine intelligence. ■

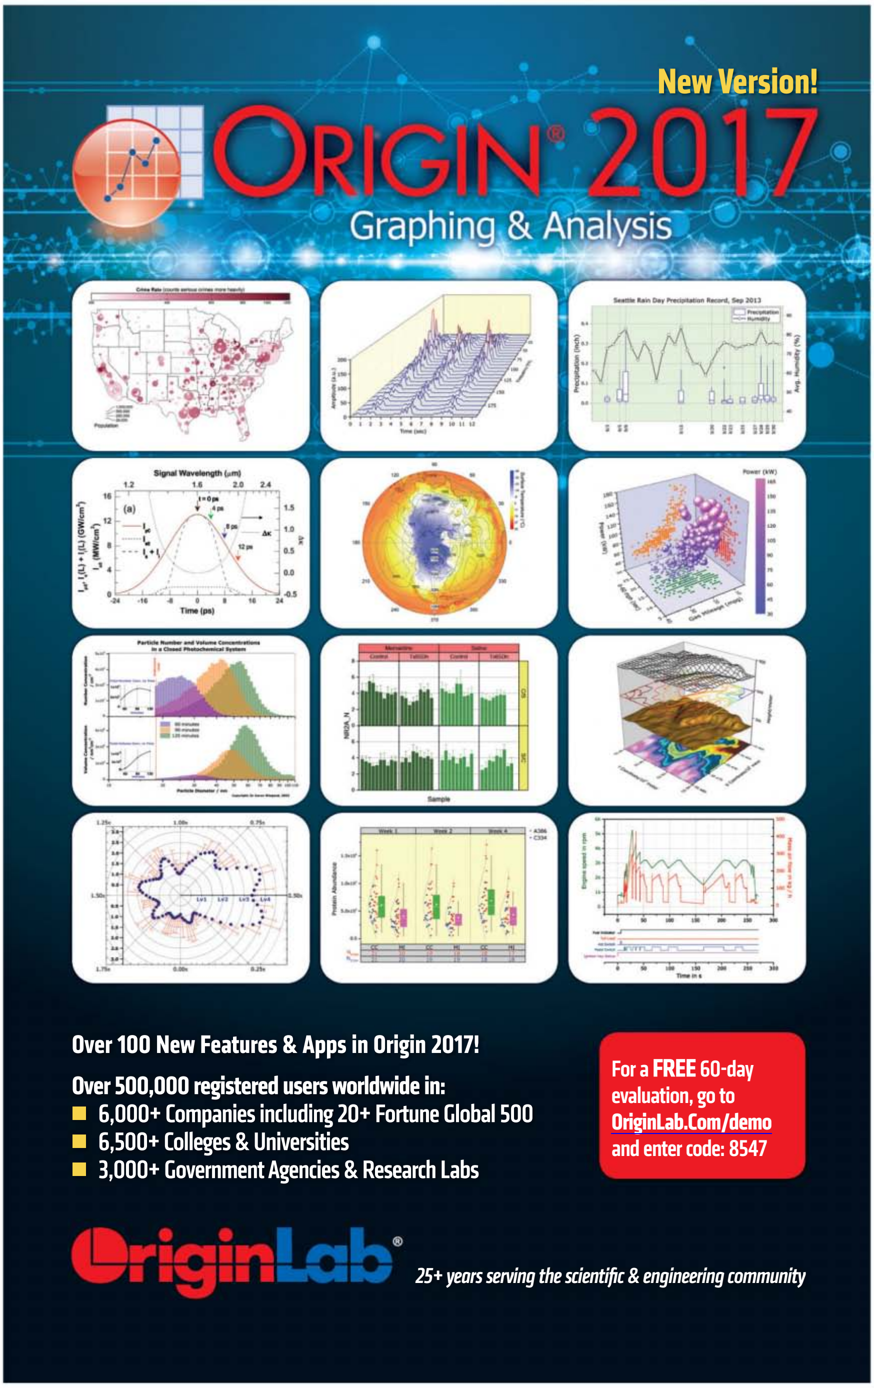

ZJU-UIUC INSTITUTE

Zhejiang University/University of Illinois at Urbana-Champaign Institute

Tenure-Track and Tenured Faculty Positions in Electrical Engineering and Computer Science

The newly launched Zhejiang University/University of Illinois at Urbana-Champaign Institute (the ZJU-UIUC Institute) is an engineering college on the new Zhejiang University (ZJU) International Campus in Haining, China, about 120 km southwest of Shanghai, invites highly qualified candidates for multiple tenure-track faculty positions at all levels and in areas of engineering and science that match its multidisciplinary mission. Candidates should have exceptional academic record or demonstrate strong potential in the cutting-edge research areas of engineering and science multidisciplinary technologies.

Duties: The Institute is built up a world-class research and education institute, it conducts teaching and research in the broad areas of engineering sciences. The Institute holds classes and student activities in English. Applications are especially encouraged from candidates whose interests address interdisciplinary topics exemplified by computer and data engineering, communications and networking, smart energy & power devices and systems, RF and THz technology, microelectronics and photonics.

Candidates must have an earned doctorate, excellent academic credentials, strong research plans, and an outstanding ability to teach effectively.

Compensation and benefits: Salary and research initiation support will be commensurate with qualifications and competitive with international norms, including housing benefits.

Applications: The search process will continue until positions are filled or through **Dec. 20, 2017**. Application materials should include a cover letter with current contact information including email address, as well as complete curriculum vitae, statements of research and teaching goals, and the names of three or more references. Please submit applications at **http://zjui.illinois.edu** or at **searchcommittee@zjui.illinois.edu**, or **zjuihr@zju.edu.cn**. For more information, please visit **http://zjui.zju.edu.cn** for job opportunities.

上海科技大学
ShanghaiTech University

TENURE-TRACK AND TENURED POSITIONS

ShanghaiTech University invites highly qualified candidates to fill multiple tenure-track/tenured faculty positions as its core founding team in the School of Information Science and Technology (SIST). We seek candidates with exceptional academic records or demonstrated strong potentials in all cutting-edge research areas of information science and technology. They must be fluent in English. English-based overseas academic training or background is highly desired. ShanghaiTech is founded as a world-class research university for training future generations of scientists, entrepreneurs, and technical leaders. Boasting a new modern campus in Zhangjiang Hightech Park of cosmopolitan Shanghai, ShanghaiTech shall trail-blaze a new education system in China. Besides establishing and maintaining a world-class research profile, faculty candidates are also expected to contribute substantially to both graduate and undergraduate educations.

Academic Disciplines: Candidates in all areas of information science and technology shall be considered. Our recruitment focus includes, but is not limited to: computer architecture, software engineering, database, computer security, VLSI, solid state and nano electronics, RF electronics, information and signal processing, networking, security, computational foundations, big data analytics, data mining, visualization, computer vision, bio-inspired computing systems, power electronics, power systems, machine and motor drive, power management IC as well as inter-disciplinary areas involving information science and technology.

Compensation and Benefits: Salary and startup funds are highly competitive, commensurate with experience and academic accomplishment. We also offer a comprehensive benefit package to employees and eligible dependents, including on-campus housing. All regular ShanghaiTech faculty members will join its new tenure-track system in accordance with international practice for progress evaluation and promotion.

Qualifications:

- Strong research productivity and demonstrated potentials;
- Ph.D. (Electrical Engineering, Computer Engineering, Computer Science, Statistics, Applied Math, or related field);
- A minimum relevant (including PhD) research experience of 4 years.

Applications: Submit (in English, PDF version) a cover letter, a 2-page research plan, a CV plus copies of 3 most significant publications, and names of three referees to: sist@shanghaitech.edu.cn. ***For more information, visit*** http://sist.shanghaitech.edu.cn/NewsDetail.asp?id=373

Deadline: The positions will be open until they are filled by appropriate candidates.

IEEE *Access*®

Multidisciplinary ⁝ Rapid Review ⁝ Open Access Journal

AMERICAN PUBLISHERS AWARDS Best New Journal in STM 2015

Become a published author in 4 to 6 weeks.

IEEE *Access* is a multidisciplinary journal that allows you to:

- Reach millions of global users through the IEEE *Xplore*® digital library with free access to all
- Submit multidisciplinary articles that do not fit neatly in traditional journals
- Expect a rapid yet rigorous peer review—a key factor why IEEE *Access* is included in Web of Science (and has an Impact Factor)
- Establish yourself as an industry pioneer by contributing to trending, interdisciplinary topics in one of the Special Sections
- Integrate multimedia and track usage and citation data for each published article
- Connect with readers through commenting
- Publish without a page limit for **only $1,750** per article

IEEE *Access*...a multidisciplinary open access journal that's worthy of the IEEE.

Learn more at: **ieeeaccess.ieee.org**

IEEE

17-PUB-013 3/17

IEEE.tv gets a mobile makeover

Bring an award-winning network of technology programs with you.

Available on the App Store

Available in Android Market

Go mobile or get the app.
www.ieee.tv

IEEE

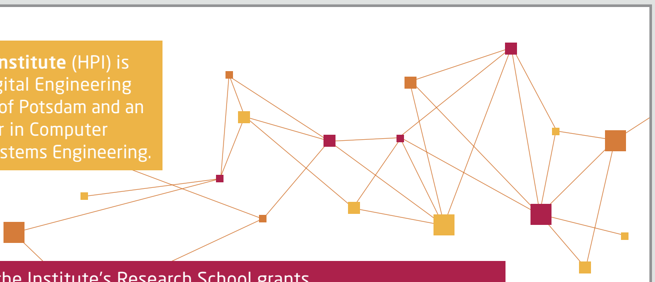

Hasso Plattner Institute (HPI) is the faculty of Digital Engineering at the University of Potsdam and an excellence center in Computer Science and IT Systems Engineering.

Annually, the Institute's Research School grants

10 Ph.D. and Postdoctoral Scholarships

With its interdisciplinary and international structure, the Research School interconnects the HPI research groups as well as its branches at University of Cape Town, Technion, and Nanjing University.

HPI RESEARCH GROUPS

- **Algorithm Engineering,** Prof. Dr. Tobias Friedrich
- **Business Process Technology,** Prof. Dr. Mathias Weske
- **Computer Graphics Systems,** Prof. Dr. Jürgen Döllner
- **Enterprise Platform and Integration Concepts,** Prof. Dr. h.c. Hasso Plattner
- **Human Computer Interaction,** Prof. Dr. Patrick Baudisch
- **Information Systems,** Prof. Dr. Felix Naumann
- **Internet Technologies and Systems,** Prof. Dr. Christoph Meinel
- **Knowledge Discovery and Data Mining,** Prof. Dr. Emmanuel Müller
- **Operating Systems and Middleware,** Prof. Dr. Andreas Polze
- **Software Architecture,** Prof. Dr. Robert Hirschfeld
- **System Analysis and Modeling,** Prof. Dr. Holger Giese

Applications must be submitted by August 15 of the respective year.
For more information on the HPI Research School please visit:
www.hpi.de/research-school

HPI Hasso Plattner Institut
Digital Engineering · Universität Potsdam

Make a world of difference

Bring the promise of technology, and the knowledge and power to leverage it, to people around the globe. **Donate now to the IEEE Foundation and make a positive impact on humanity.**

- **Inspire technology education**
- **Enable innovative solutions for social impact**
- **Preserve the heritage of technology**
- **Recognize engineering excellence**

IEEE Foundation

Discover how you can contribute today.

Learn more about the IEEE Foundation at **ieeefoundation.org.** To make a donation now, go to **ieeefoundation.org/donate.**

IEEE

SFU SCHOOL OF ENGINEERING SCIENCE

Simon Fraser University School of Engineering Science Tenure-track Faculty Position

The School of Engineering Science at Simon Fraser University, British Columbia, Canada, invites applications for a tenure-track faculty position in Computer Engineering, mainly at the Assistant Professor level, starting as early as Fall 2017. Outstanding candidates at Associate and Full Professor levels may also be considered. Exceptional candidates may also be considered for a Canada Research Chair at an appropriate tier. Strategically important research expertise that builds on or complements existing strengths is sought.

Candidates at the Assistant Professor level are expected to demonstrate a commitment to excellence in research, graduate student supervision, and teaching in computer engineering at the undergraduate and graduate levels. Senior applicants will have already demonstrated such excellence. The ability to bridge disciplines is desirable. A Ph.D. in Electrical and Computer Engineering or related areas by the date of appointment is required. Candidates should have P.Eng. accreditation or at a minimum secure such qualification as soon as possible after appointment.

Simon Fraser University is located in metropolitan Vancouver, one of the most livable cities in the world and has been ranked by Maclean's Magazine as the top comprehensive university in Canada for 14 times in the last 25 years. The School of Engineering Science currently has 29 faculty members, 120 graduate students, and more than 1000 undergraduate students. The school prides itself in its excellent researchers and facilities, as well as high-quality engaged students. The School of Engineering Science is home to the current Natural Sciences and Engineering Research Council (NSERC) Chair for Women in Science and Engineering.

The review of applications will begin May 29, 2017 and will continue until the position is filled. The position is subject to availability of funding and approval by the SFU Board of Governors. To apply, please submit your curriculum vitae, research and teaching statement, plus the names and email addresses of three referees to our online application system located at **https://confs.precisionconference.com/~ensc17a/apply**

All qualified candidates are encouraged to apply. However, Canadian citizens and permanent residents will be given priority. Simon Fraser University is committed to employment equity and encourages applications from all qualified women and men, including visible minorities, aboriginal people, and persons with disabilities. Under the authority of the University Act, personal information that is required by the University for academic appointment competitions will be collected. For further details see the Collection Notice: **http://www.sfu.ca/vpacademic/faculty_openings/collection_notice.html**

PAST FORWARD_BY EVAN ACKERMAN

A ROBOT'S DELICATE TOUCH

The Mobot Mark II, developed by Hughes Aircraft Co. in 1961, was a teleoperated mobile manipulator. A cable connected it to a remote-control panel, from which a human operator directed the robot's motions while viewing the scene through a pair of TV cameras. The Mobot was designed to be used for situations deemed far too dangerous for humans, like handling radioactive materials, exploring the moon, and helping women get dressed. ■

J. R. EYERMAN/THE LIFE PICTURE COLLECTION/GETTY IMAGES