FOR THE TECHNOLOGY INSIDER
FEBRUARY 2023

IEEE Spectrum

INNER ELECTRONICS

THE HIDDEN BEAUTY OF THE SIMPLEST DEVICES

IEEE

VOLUME 60 / ISSUE 2 FEBRUARY 2023

Contents

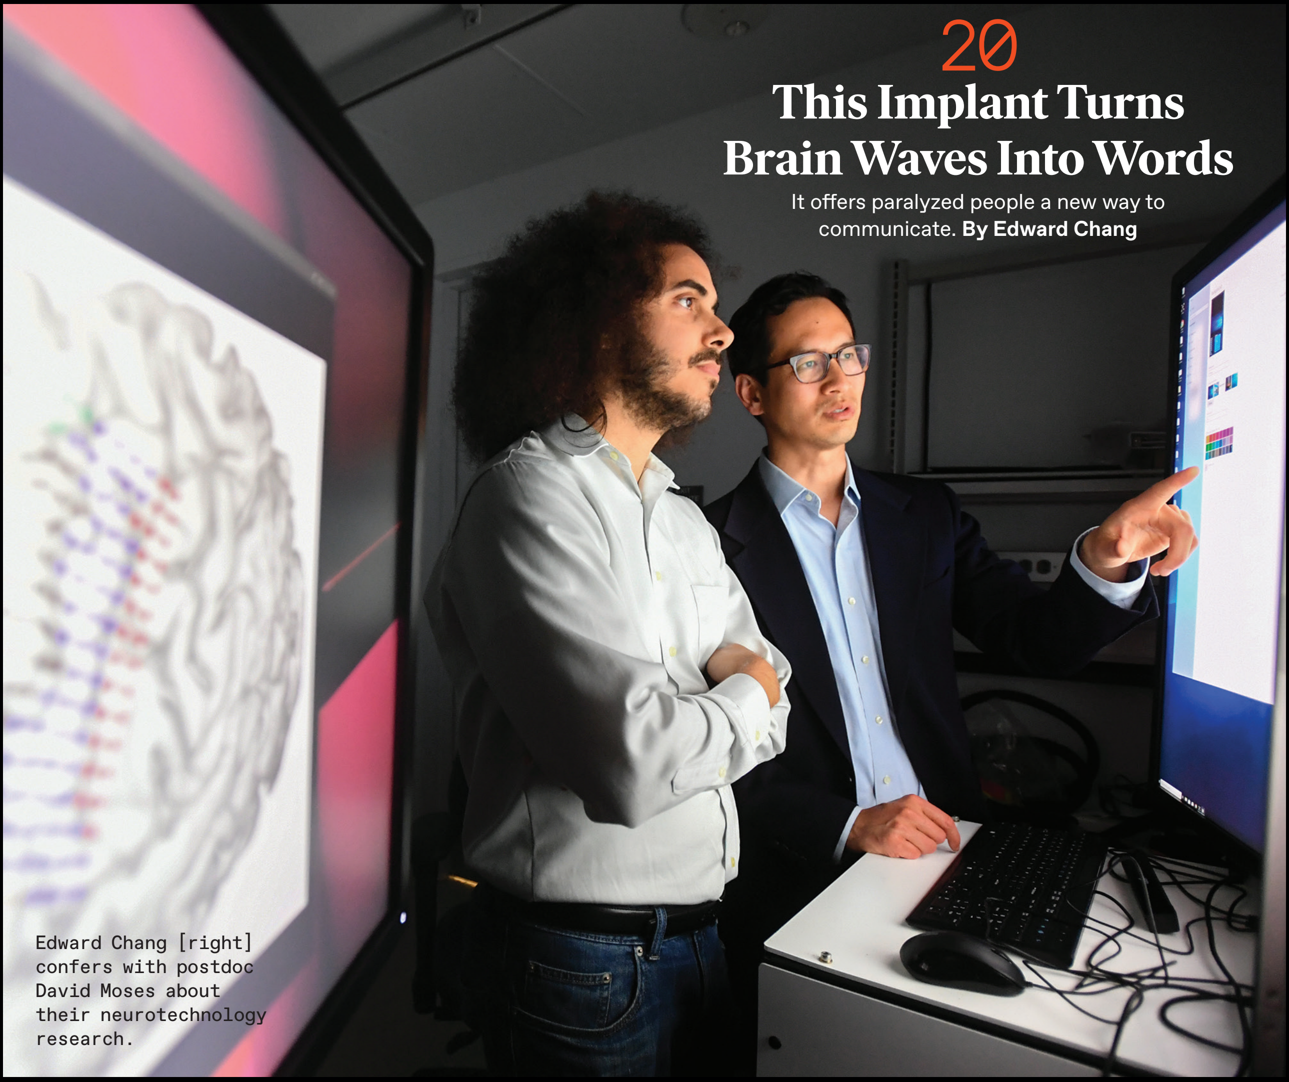

Edward Chang [right] confers with postdoc David Moses about their neurotechnology research.

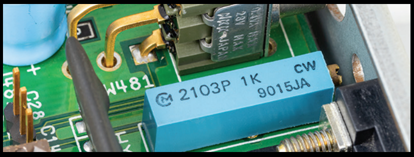

ON THE COVER:
Photograph by Eric Schlaepfer & Windell H. Oskay

TOP: NOAH BERGER; BOTTOM: ERIC SCHLAEPFER & WINDELL H. OSKAY

BY HARRY GOLDSTEIN

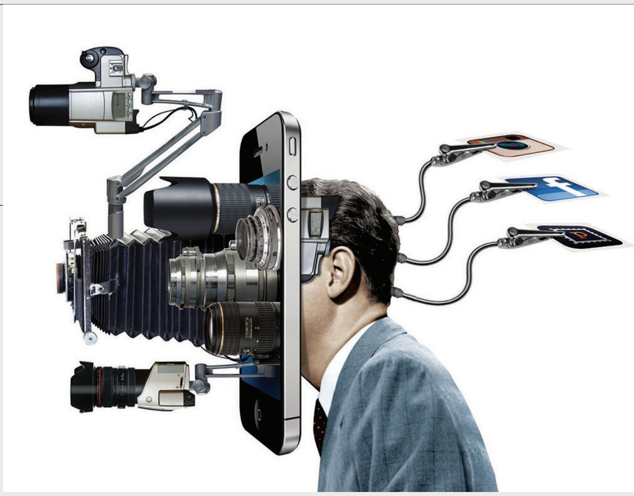

Transparency Depends on Digital Breadcrumbs

But is transparency alone sufficient for a free and open society?

From the moment we wake up and reach for our smartphones and throughout the day as we text each other, upload selfies to social media, shop, commute, work, work out, watch streaming media, pay bills, and travel, and even while we're sleeping, we spew personal data like jets sketching contrails across the sky.

An astonishing amount of that data is recorded, stored, analyzed, and shared by media companies looking to pitch you content and ads, retailers aiming to sell you more of what you've already bought, and potential distant relatives hitting you up on genealogy sites. And sometimes, if you're suspected of participating in illegal activities, that data can bring you under the scrutiny of law enforcement officials.

Contributing Editor Mark Harris spent months poring over court documents and other records to understand how the U.S. Federal Bureau of Investigation and other agencies exploited vast troves of data to conduct the largest criminal investigation in U.S. history: into the violent overtaking of the Capitol building on 6 January 2021.

The events of that day unfolded on live television watched by millions. But in order for investigators to identify suspects amid a mob of thousands, they had to cast a very wide net and sought the cooperation of tech giants like Google, Facebook, and Snap and carriers like Verizon and T-Mobile. As Harris painstakingly documents in "The Panopticon v. the Capitol Rioters" [p. 32], some of the information the FBI used was intentionally shared by rioters on social media, while other information was gleaned from the kind of data we all heedlessly cast off during the course of the day, like the order for pizza that landed one group of rioters in hot water or the automated license-plate readings that were cited in 20 cases.

> **"In the eternal struggle between security and privacy, the best that digital-rights activists can hope for is to watch the investigators as closely as they are watching us."**
> —MARK HARRIS

The ability to ingest multiple data streams and analyze them to trace rioters' journeys to, through, and back from the Capitol has led to 950 arrests, with more than half leading to guilty pleas and 40 to guilty verdicts as of this writing. But as the privacy advocates Harris interviewed point out, while these tools helped law enforcement hold some people accountable for their actions that day, those same tools can be used by the state against law-abiding citizens, not just in the United States, but anywhere. And the data we make available (knowingly or not), often for the sake of convenience or as the price of admission, leaves us vulnerable to bad actors, be they governments, corporations, or individuals.

The writer David Brin foretold a version of our current panopticon in his 1998 book *The Transparent Society*. In it, he acknowledges the risks of surveillance technology but contends that the very ubiquity of that technology is in itself a safeguard against abuse by giving everyone the ability to shine a light on the dark corners of individual and institutional behavior. His stance jibes with Harris's final observation: "In the eternal struggle between security and privacy, the best that digital-rights activists can hope for is to watch the investigators as closely as they are watching us."

But as Brin points out, watching the watchers isn't enough to guarantee a free and open society. As data-driven prosecutions for the 6 January insurrection continue, it's worth considering the linkage Brin makes between liberty and accountability, which he says, "is the one fundamental ingredient on which liberty thrives. Without the accountability that derives from openness—enforceable upon even the mightiest individuals and institutions—how can freedom survive?" ■

PORTRAIT BY SERGIO ALBIAC; ILLUSTRATION BY VIKTOR KOEN

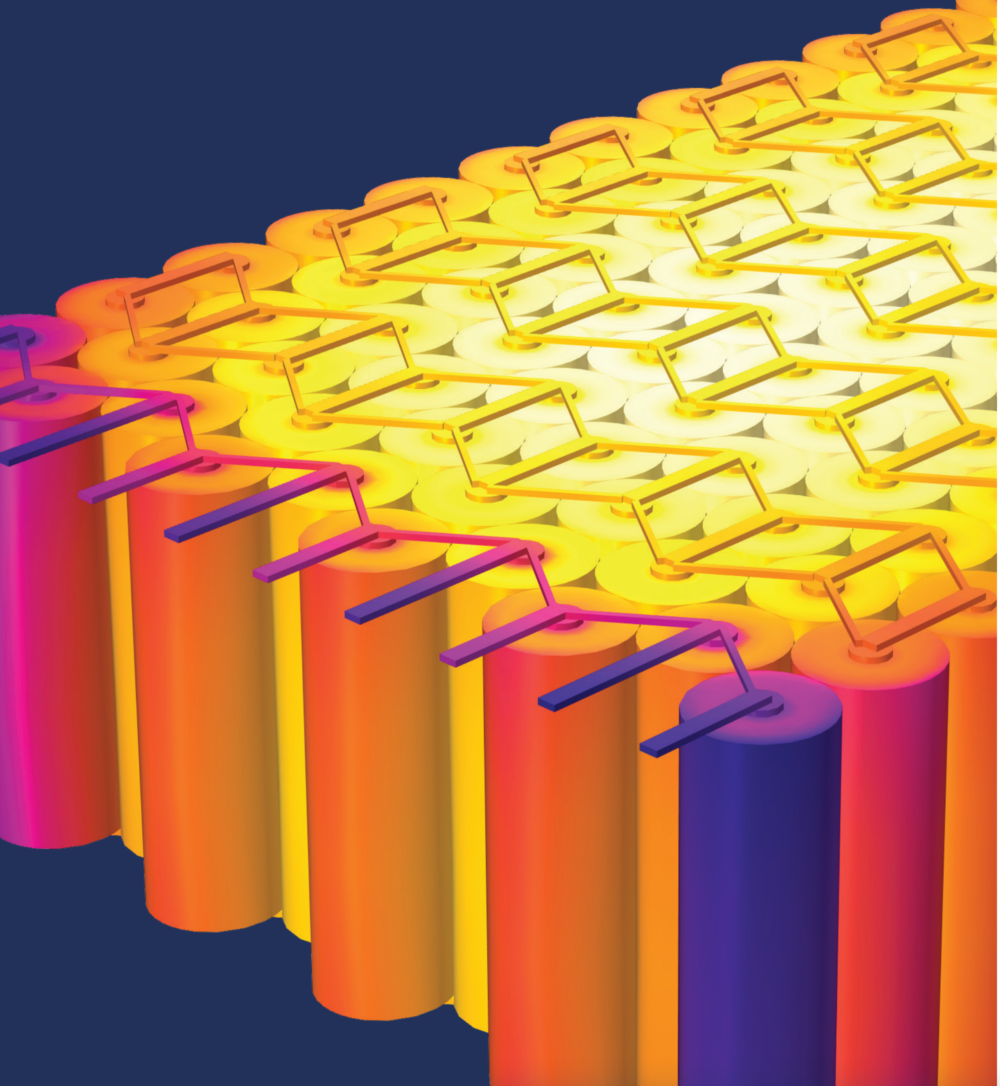
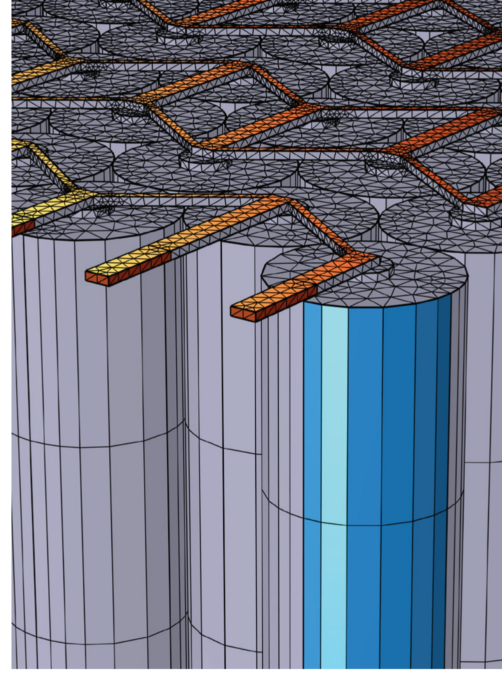
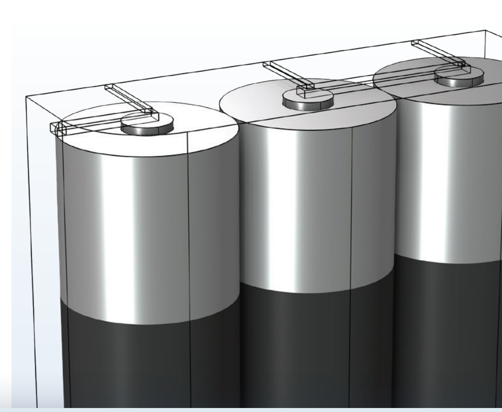

CONTRIBUTORS

EDWARD CHANG

In this issue, Chang, the chair of neurological surgery at the University of California, San Francisco writes about developing technology for patients of his who have lost the ability to speak [p. 20]. His lab works on decoding brain signals associated with intended speech, a project that requires not only today's best neurotechnology hardware but also powerful machine-learning models. "We were on a collision course with what's going on in artificial intelligence," says Chang. "That's what's enabled all of this."

MARK HARRIS

Harris, an investigative-technology reporter based in Seattle, dove deeply into the digital records that police amassed on the 6 January rioters for his story [p. 32]. "I've personally read and categorized well over 1,000 charging documents—weeks of work," he told us, when he filed his draft. Looking at the scale of the police effort, Harris asks a simple question: Can we, the public, watch the investigators as closely as they are watching us?

ERIC SCHLAEPFER

Schlaepfer and Windell H. Oskay are coauthors of *Open Circuits: The Inner Beauty of Electronic Components* (No Starch Press, 2022), which we excerpt in this issue [p. 38]. For over a decade, the pair have been collaborating on projects that require some "extraordinary engineering but also a sort of artistic flourish," says Oskay. One of their current collaborations is a 6502 microprocessor built entirely out of discrete transistors. While there are several working prototypes, a commercial product is held up by "a logjam in the supply chain," he says.

HALLAM STEVENS

Stevens is a professor of interdisciplinary studies at James Cook University, in Townsville, Australia. He currently researches the history of information technology and biotechnology in Asia. For a decade, Stevens lived in Singapore, home of Henn Tan, the now-jailed inventor of the USB memory stick, whose story he tells on page 26. Stevens is the proud owner of one of Tan's early 16-megabyte Trek ThumbDrives.

IEEE Spectrum

ACTING EDITOR IN CHIEF Harry Goldstein, h.goldstein@ieee.org

EXECUTIVE EDITOR Jean Kumagai, j.kumagai@ieee.org

MANAGING EDITOR Elizabeth A. Bretz, e.bretz@ieee.org

CREATIVE DIRECTOR
Mark Montgomery, m.montgomery@ieee.org

PRODUCT MANAGER, DIGITAL Erico Guizzo, e.guizzo@ieee.org

EDITORIAL DIRECTOR, CONTENT DEVELOPMENT
Glenn Zorpette, g.zorpette@ieee.org

SENIOR EDITORS
Evan Ackerman (Digital), ackerman.e@ieee.org
Stephen Cass (Special Projects), cass.s@ieee.org
Samuel K. Moore, 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
Eliza Strickland, e.strickland@ieee.org

ART & PRODUCTION
DEPUTY ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ONLINE ART DIRECTOR Erik Vrielink, e.vrielink@ieee.org
PRINT PRODUCTION SPECIALIST
Sylvana Meneses, s.meneses@ieee.org
MULTIMEDIA PRODUCTION SPECIALIST
Michael Spector, m.spector@ieee.org

NEWS MANAGER Margo Anderson, m.k.anderson@ieee.org

ASSOCIATE EDITORS
Willie D. Jones (Digital), w.jones@ieee.org
Michael Koziol, m.koziol@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
EDITORIAL INTERN Dina Genkina, d.genkina@ieee.org

CONTRACT SPECIALIST Ramona L. Foster, r.foster@ieee.org

CONTRIBUTING EDITORS Robert N. Charette, Steven Cherry, Charles Q. Choi, Peter Fairley, Edd Gent, W. Wayt Gibbs, Mark Harris, Allison Marsh, Prachi Patel, Julianne Pepitone, Lawrence Ulrich, Emily Waltz

THE INSTITUTE
EDITOR IN CHIEF Kathy Pretz, k.pretz@ieee.org
ASSOCIATE EDITOR Joanna Goodrich, j.goodrich@ieee.org

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy

ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org
SENIOR ADVERTISING PRODUCTION COORDINATOR
Nicole Evans Gyimah, n.gyimah@ieee.org
ADVERTISING PRODUCTION +1 732 562 6334

EDITOR IN CHIEF EMERITUS Susan Hassler, s.hassler@ieee.org

EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM*
Susan Hassler, *Chair*; Robert Schober, Ella M. Atkins, Sangyeun Cho, Francis J. "Frank" Doyle III, Hugh Durrant-Whyte, Matthew Eisler, Shahin Farshchi, Alissa Fitzgerald, Benjamin Gross, Lawrence O. Hall, Daniel Hissel, Jason K. Hui, Michel M. Maharbiz, Somdeb Majumdar, Lisa May, Carmen S. Menoni, Ramune Nagisetty, Paul Nielsen, Sofia Olhede, Christopher Stiller, Mini S. Thomas, Wen Tong, Haifeng Wang, Boon-Lock Yeo

EDITORIAL ADVISORY BOARD, *THE INSTITUTE*
Kathy Pretz, *Chair*; Qusi Alqarqaz, Stamatis Dragoumanos, Jonathan Garibaldi, Madeleine Glick, Lawrence O. Hall, Susan Hassler, Francesca Iacopi, Cecilia Metra, Shashi Raj Pandey, John Purvis, Chenyang Xu

MANAGING DIRECTOR, PUBLICATIONS Steven Heffner

DIRECTOR, BUSINESS DEVELOPMENT, MEDIA & ADVERTISING Mark David, m.david@ieee.org

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

ADVERTISING INQUIRIES Naylor Association Solutions,
Erik Albin +1 352 333 3371, ealbin@naylor.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*.

IEEE

IEEE BOARD OF DIRECTORS
PRESIDENT & CEO Saifur Rahman, president@ieee.org
+1 732 562 3928 Fax: +1 732 981 9515
PRESIDENT-ELECT Thomas M. Coughlin
TREASURER Mary Ellen Randall
SECRETARY Forrest "Don" Wright
PAST PRESIDENT K.J. Ray Liu

VICE PRESIDENTS
Rabab Ward, *Educational Activities*; Sergio Benedetto, *Publication Services & Products*; Jill I. Gostin, *Member & Geographic Activities*; John P. Verboncoeur, *Technical Activities*; Yu Yuan, *President, Standards Association*; Eduardo F. Palacio, *President, IEEE-USA*

DIVISION DIRECTORS
Franco Maloberti (I); Kevin L. Peterson (II); Khaled Ben Letaief (III); Alistair P. Duffy (IV); Cecilia Metra (V); Kamal Al-Haddad (VI); Claudio Cañizares (VII); Leila De Floriani (VIII); Ali H. Sayed (IX); Stephanie M. White (X)

REGION DIRECTORS
Greg T. Gdowski (1); Andrew D. Lowery (2); Theresa A. Brunasso (3); Vickie A. Ozburn (4); Bob G. Becnel (5); Kathy Hayashi (6); Robert L. Anderson (7); Vincenzo Piuri (8); Enrique A. Tejera (9); ChunChe (Lance) Fung (10)

DIRECTOR EMERITUS Theodore W. Hissey

IEEE STAFF
EXECUTIVE DIRECTOR & COO Sophia A. Muirhead
+1 732 562 5400, s.muirhead@ieee.org
CHIEF INFORMATION OFFICER Cherif Amirat
+1 732 562 6017, c.amirat@ieee.org
CHIEF HUMAN RESOURCES OFFICER Liesel Bell
+1 732 562 6347, l.bell@ieee.org
CHIEF MARKETING OFFICER Karen L. Hawkins
+1 732 562 3964, k.hawkins@ieee.org
PUBLICATIONS Steven Heffner
+1 212 705 8958, s.heffner@ieee.org
CORPORATE ACTIVITIES Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MEMBER & GEOGRAPHIC ACTIVITIES Cecelia Jankowski
+1 732 562 5504, c.jankowski@ieee.org
STANDARDS ACTIVITIES Konstantinos Karachalios
+1 732 562 3820, constantin@ieee.org
EDUCATIONAL ACTIVITIES Jamie Moesch
+1 732 562 5514, j.moesch@ieee.org
CHIEF FINANCIAL OFFICER Thomas R. Siegert
+1 732 562 6843, t.siegert@ieee.org
TECHNICAL ACTIVITIES Mary Ward-Callan
+1 732 562 3850, m.ward-callan@ieee.org
ACTING MANAGING DIRECTOR, IEEE-USA Russell T. Harrison
+1 202 530 8326, r.t.harrison@ieee.org

IEEE PUBLICATION SERVICES & PRODUCTS BOARD
Sergio Benedetto, *Chair*; Stefano Galli, Maria Sabrina Greco, Lawrence O. Hall, James Irvine, Charles M. Jackson, Clem Karl, Yong Lian, Fabrizio Lombardi, Aleksandar Mastilovic, Anna Scaglione, Gaurav Sharma, Isabel Trancoso, Peter Winzer, Weihua Zhuang

IEEE OPERATIONS CENTER
445 Hoes Lane, Box 1331
Piscataway, NJ 08854-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2023 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 60, Issue No. 2. 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, ℅ 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 https://www.ieee.org/about/corporate/governance/p9-26.html.

News

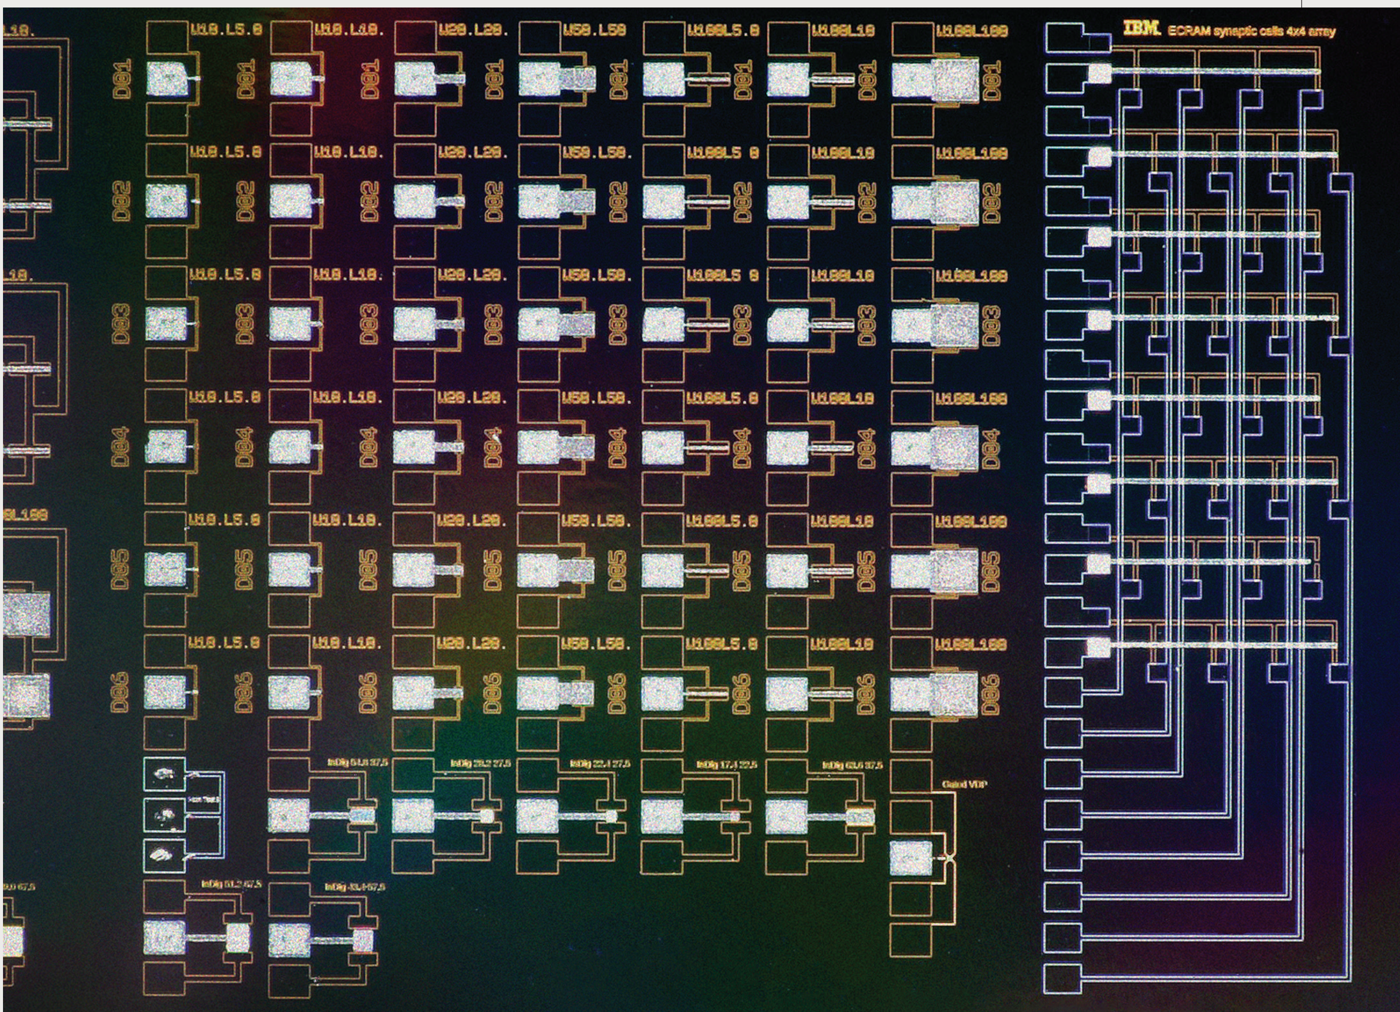

This analog electrochemical memory (ECRAM) array provides a prototype for artificial synapses in AI training.

IBM RESEARCH

ARTIFICIAL INTELLIGENCE

Better AI Through Chemistry › Battery-inspired device enables fast analog AI

BY DINA GENKINA

How far away could an artificial brain be? Perhaps a very long way off still, but a working analogue to the essential element of the brain's networks, the synapse, appears closer at hand now.

That's because a device that draws inspiration from batteries now appears surprisingly well suited to run artificial neural networks. Called electrochemical RAM (ECRAM), it is giving traditional transistor-based AI an unexpected run for its money—and is quickly moving toward the head of the pack in the race to develop the perfect artificial synapse. Researchers reported a string of advances in late 2022 at the IEEE International Electron Device Meeting (IEDM 2022) and elsewhere, including ECRAM devices that use less energy, hold memory longer, and take up less space than their predecessors.

The artificial neural networks that power today's machine-learning algorithms are based on software that models a large collection of electronics-based "neurons," along with their many connections, or synapses. Instead of representing neural networks in software, researchers think that faster, more energy-efficient AI would result from representing the components, especially the synapses, with real devices. This concept, called analog AI, requires a memory cell that combines a whole slew of difficult-to-obtain properties: It needs to have a broad enough range of analog values, switch between those values reliably and quickly, while also holding onto those values for a long time, while *also* still being amenable to manufacturing at scale.

Most types of memory are well adapted to store digital values but are too noisy to reliably store analog. But back in 2016, a group of researchers at Sandia National Laboratories led by Alec Talin realized that the answer was right in front of them: the state of charge of a battery. "Fundamentally, a battery works by moving ions between two materials. As the ion moves between the two materials, the battery stores and releases energy," says Yiyang Li, then at Sandia but now an assistant professor of materials science and engineering at the University of Michigan. "We found that we can use the same process for storing information."

In other words, the number of ions in the channel determines the stored analog value. Theoretically, a difference of a single ion could be detectable. ECRAM uses these concepts by controlling how much charge is in the "battery" through a third gate terminal.

"These devices responded much faster than the brain synapse."

—JESUS DEL ALAMO, MIT

Such a battery would have a negative terminal, an ion-doped channel in the middle, and a positive terminal on the other end. The conductivity between the positive and negative terminal, prescribed by the number of ions in the channel, is what determines the analog value stored in the device. Above the channel, there's an electrolyte barrier that permits ions (but not electrons) through. On top of the barrier is a reservoir layer, containing a supply of mobile ions. A voltage applied to this reservoir serves as a "gate," forcing ions through the electrolyte barrier into the channel, or the reverse. These days, the time it takes to switch to any desired stored value is phenomenally fast.

"These devices responded much faster than the brain synapse," says Jesus del Alamo, professor of engineering and computer science at MIT. "As a result, they give us the possibility of essentially being able to do a brainlike computation, artificial-intelligence computation, significantly faster than the brain, which is what we really need to realize the promise of artificial intelligence."

Del Alamo's group at MIT has opted for individual protons as the device's primary information carrier, because of their unparalleled speed compared to larger ions. Just a few months ago, the researchers demonstrated devices that move ions around in mere nanoseconds, about 10,000 times as fast as synapses in the brain. But fast was not enough.

"We can see the device responding very fast to [voltage] pulses that are still a little bit too big," del Alamo says, "and that's a problem. We want to be able to also get the devices to respond very fast with pulses that are of lower voltage because that is the key to energy efficiency."

In research reported in December at IEDM 2022, the MIT group dug down into the details of their device's operation with the first real-time study of current flow. They discovered what they believe is a bottleneck that prevents the devices from switching at lower voltages: The protons traveled easily across the electrolyte layer but needed an extra voltage push at the interface between the electrolyte and the channel. Armed with this knowledge, researchers believe they can engineer the material interface to reduce the voltage required for switching, opening the door to higher energy efficiency and scalability, says del Alamo.

Then, once programmed, these devices usually hold resistivity for a few hours. Researchers at Sandia and the University of Michigan have teamed up to push the envelope on this retention time—to 10 years. They published their results in the journal *Advanced Electronic Materials* in November.

To retain memory for this long, the team, led by Michigan's Li, opted for the heavier oxygen ion instead of the proton in the MIT device. Even with a more massive ion, what they observed was unexpected. "I remember one day, while I was traveling, my graduate student Diana Kim showed me the data, and I was astounded, thinking something was incorrectly done," recalls Li. "We did not expect it to be so nonvolatile. We later repeated this over and over, before we gained enough confidence."

They speculate that the nonvolatility comes from their choice of material, tungsten oxide, and the way oxygen ions arrange themselves inside it. "We think it's due to a material property called phase separation that allows the ions to arrange themselves such that there's no driving force pushing them back," Li explains.

Unfortunately, this long retention time comes at the expense of switching speed, which is measured in minutes for Li's device. But, the researchers say, having a physical understanding of how the retention time is achieved enables them to look for other materials that

Nanosecond programming time, years-long storage time—no single technology can yet do both.

show a long memory and faster switching properties simultaneously.

Meanwhile, the fact that these devices have an added third terminal makes them bulkier than competing two-terminal memories, which could limit scalability.

So, to help shrink the devices and pack them efficiently into an array, researchers at Pohang University of Science and Technology, in South Korea, laid them on their side. This allowed the researchers to reduce the devices to a mere 30-by-30-nanometer footprint, an area about one-fifth as large as that of previous generations, while retaining switching speed and even improving on the energy efficiency and read time. They also reported their results at IEDM 2022.

The team structured their device as one big vertical stack: The source was deposited on the bottom, the conducting channel was placed next, then the drain above it. To allow the drain to permit ions into and out of the channel, the researchers replaced the usual semiconductor material with a single layer of graphene. This graphene drain also served as an extra barrier controlling the ion flow. Above it, they placed the electrolyte barrier, and finally the ion reservoir and gate terminal on top. With this configuration, not only did the performance not degrade, but the energy required to write and read information into the device decreased. And, as a result, the time required to read the stored analog value fell by a factor of 20.

Even with all the above advances, a commercial ECRAM chip that accelerates AI training is still some distance away. The devices can now be made of foundry-friendly materials, but that's only part of the story, says John Rozen, program director at the IBM Research AI Hardware Center. "A critical focus of the community should be to address integration issues to enable ECRAM devices to be coupled with front-end transistor logic monolithically on the same wafer, so that we can build demonstrators at scale and establish if it is indeed a viable technology."

Rozen's team at IBM is working toward this manufacturability. In the meantime, they've created a software tool that allows the user to play around with using different emulated analog AI devices, including ECRAM, to actually train neural networks and evaluate their performance. ■

AEROSPACE

Moon Rover's Test Run

Part of the preparation for the future moon missions to take place under the heading of the Artemis project is testing equipment such as this moon rover prototype on the volcanic soil of Black Point Lava Flow near Flagstaff, Ariz.

OLIVIER TOURON/AFP/GETTY IMAGES

MARINA PAPADOPOULOU/UNIVERSITY OF GRONINGEN

ROBOTS

Robotic Falcon Is the Scarecrow of the Skies › The raptorlike drone keeps birds away from airports

BY EDD GENT

Collisions with birds cost the civil aviation industry roughly US $1.3 billion and kill thousands of animals every year. Perhaps a robotic imitation of a peregrine falcon could scare birds away from airports, where most such collisions occur.

Airports have often tried to scare birds away with loud pyrotechnics or speakers that play avian distress calls. But these approaches tend to become ineffective as the birds get desensitized by repeated exposure, says Charlotte Hemelrijk, a professor on the faculty of science and engineering at the University of Groningen, in the Netherlands. Live hawks or blinding lasers are also sometimes used to disperse flocks, she says, but this is controversial as it can harm the animals. Besides, keeping and training falcons is not cheap.

So Hemelrijk and colleagues designed a flying robot the same size and shape as a hawk and painted its fiberglass and carbon-fiber body to mimic the markings of a real falcon.

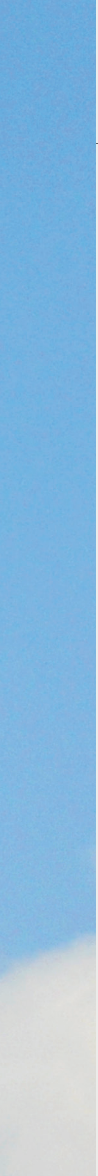

Dutch researchers have developed a big, hawk-shaped drone [left] to chase birds from the vicinity of airports so they are not a danger to aircraft.

ROLF F. STORMS/UNIVERSITY OF GRONINGEN

The RobotFalcon doesn't flap. Instead, it relies on two small battery-powered propellers that fly it at around 48 kilometers per hour (30 miles per hour) for up to 15 minutes. A remote human operator gets a hawk's-eye perspective via a camera perched above the robot's head.

The researchers field-tested it against a conventional quadcopter drone near the Dutch city of Workum. They also compared their results to 15 years of data collected by the Royal Netherlands Air Force that assessed the effectiveness of conventional deterrence methods, such as pyrotechnics and distress calls.

In a paper published in the *Journal of the Royal Society Interface*, the team showed that the RobotFalcon cleared birds from fields faster and more effectively than the drone. It also kept birds away from fields longer than distress calls, the most effective of the conventional approaches.

The birds did not get habituated to the RobotFalcon over three months of testing, says Hemelrijk. They were also much more likely to behave as if they were trying to escape from a predator than they did when facing the conventional drone. "The birds don't distinguish it from a real falcon, it seems," says Hemelrijk.

Other attempts to use hawk-imitating robots to disperse birds have had less promising results, though. Morgan Drabik-Hamshare, a research wildlife biologist at the National Wildlife Research Center, which is a division of the U.S. Department of Agriculture, and her colleagues published a paper in *Scientific Reports* last year that described how they pitted a robotic peregrine falcon with flapping wings against a quadcopter and a fixed-wing remote-controlled aircraft.

They found the robotic falcon was the least effective of the three at scaring away turkey vultures, with the quadcopter scaring off the most birds and the remote-controlled plane eliciting the quickest response. "Despite the predator silhouette, the vultures did not perceive the predator UAS [unmanned aircraft system] as a threat," Drabik-Hamshare wrote in an email.

Zihao Wang, an associate lecturer at the University of Sydney, who develops UAS for bird deterrence, says the RobotFalcon does seem to be effective at dispersing flocks. But he points out that its wingspan is nearly twice the diagonal length of the quadcopter it was compared with, which means it creates a much larger silhouette when viewed from the birds' perspective. This means the birds could be reacting more to its size than its shape, so he would like to see the RobotFalcon compared with a similar size drone in the future.

The unique design also means the robot requires an experienced and specially trained operator, Wang adds, which could make it difficult to roll out widely. A potential solution could be to make the system autonomous, he says, but this is not necessarily easy.

Hemelrijk says automating the RobotFalcon is probably not feasible, both due to strict limits on the use of autonomous drones near airports as well as the sheer technical complexity. The robot's current operator is a falconer with significant experience, she says. Creating an autonomous system that could recognize and target bird flocks in a similar way would be highly challenging.

But Hemelrijk points out that most airports already have full-time staff dedicated to bird deterrence, and that they could be trained to use the new technology. And given the apparent lack of habituation and the ability to chase birds in a specific direction—away from runways—she thinks the robotic falcon could be a useful new spin on the old, time-tested scarecrow idea. ■

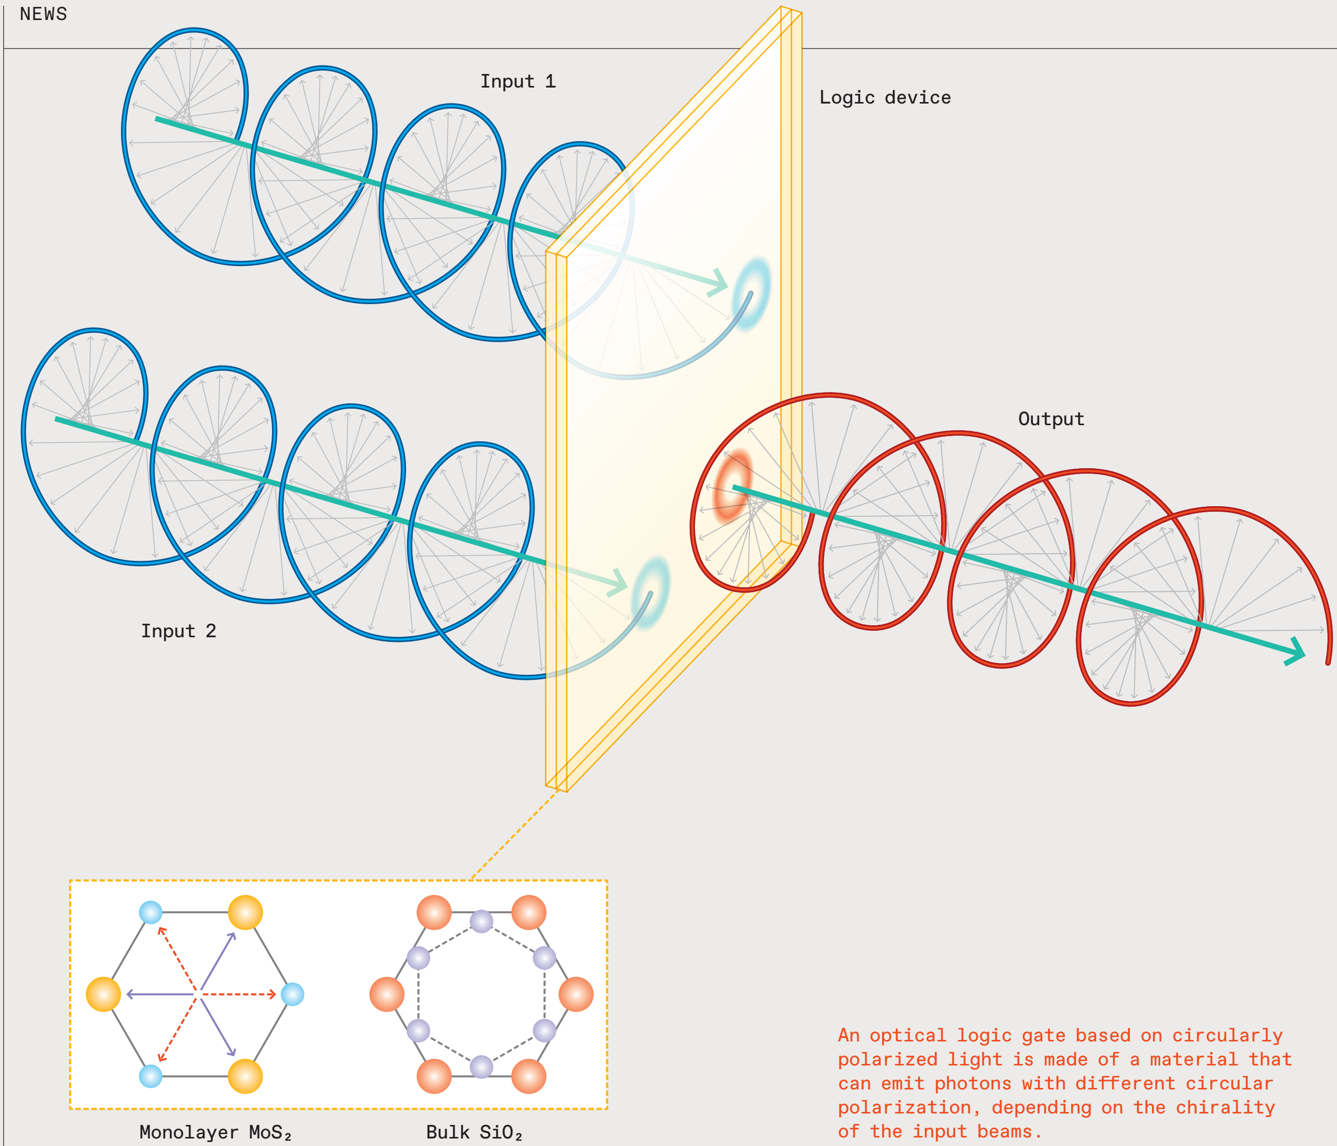

An optical logic gate based on circularly polarized light is made of a material that can emit photons with different circular polarization, depending on the chirality of the input beams.

COMPUTING

These Optical Gates Provide Electronic Access › Ultrafast optical computing can now interface with traditional circuits

BY CHARLES Q. CHOI

Logic gates built from transistors carry out operations such as AND, OR, and NOT. In recent decades, scientists have been trying to build electronic gates' optical equivalents. A new study, published on 9 December in the journal *Science Advances*, brings such optical computing closer while also revealing a promising way to connect it to conventional, electronic circuits.

The key is a property of light known as chirality, says lead author Yi Zhang, a researcher at Aalto University, in Finland. As a ray of light moves forward, it can be made to spiral much like threads on a screw, turning either clockwise or counterclockwise, in a right- or left-handed circular polarization.

Illustration by Chris Philpot

To create the gate, Zhang and his colleagues took a single layer of molybdenum disulfide—a sheet of molybdenum atoms sandwiched between two layers of sulfur atoms—and placed it on top of silicon dioxide. Then the researchers shined two light beams at the gate. When both input beams had the same chirality, or handedness, the output was right-handed, but when both input beams had different chirality, the output beam was left-handed.

This new device served as one kind of logic gate, XNOR. By adding filters or other optical components, the researchers created the AND, OR, NOR, XOR, and NAND gates. The gates switched in less than 100 femtoseconds, roughly one million times as fast as electronic gates. Moreover, the scientists found they could achieve high-speed electric control of the gates simply by applying a voltage to the molybdenum disulfide.

"Traditionally, the connections between electronic and optical computing have mainly been realized through slow and inefficient optical-to-electrical and electrical-to-optical conversion," Zhang says. "We demonstrate electrical control of the chirality optical gates, realizing an exciting prospect for direct interconnection between electrical and optical computing."

In addition, the researchers showed that a single device could simultaneously run multiple gates. In contrast, previous electronic and optical gates each typically performed just one logic operation at a time, Zhang notes. These findings suggest that simultaneous multiple-chirality logic gates could help build complex multifunctional circuits and networks, he says.

In the future, the researchers want to link the gates in a large-scale circuit so that the operations of the components can "cascade." Although previous optical gates have faced major difficulties while cascading, Zhang suggests that it theoretically should not be a problem with his team's devices.

Zhang notes that the biggest challenge they face is the very low efficiency of the nonlinear optical effect underlying their gate's operation. "The good news is that there are several new materials reported recently that have high nonlinear conversion efficiency," he says. ■

JOURNAL WATCH

Millimeter Waves Make This ECG Contactless

To measure the electrical activity of the heart, an electrocardiogram (ECG) must connect to the skin through a web of electrodes. Even the latest Apple Watch requires the user to touch the device's protruding "digital crown" with a finger to complete a circuit across the user's body.

But now researchers in China report a novel ECG technology that uses millimeter-wave radar and AI to infer an ECG signal, making the system completely contactless. If this work bears out, the technology could provide a reliable and uninterrupted stream of heart-health data.

"The need to attach electrodes to the body during current ECG monitoring decreases people's willingness to wear such devices for a long time, making those transient irregular ECG signals hard to detect," explains Yan Chen, professor in the School of Cyberspace Security at the University of Science and Technology of China.

Chen experienced this firsthand when he needed ECG monitoring for 24 hours. "During that time, I was suffering from skin irritation where the electrodes were placed and was annoyed by my limited ability to move due to the electrode wires. This experience really makes me refuse to do another examination," he says.

In their study, the researchers conducted 200 experimental trials involving 35 participants between the ages of 18 and 65. The radar device was placed between 0.4 and 0.5 meters above their bodies during four physiological states—normal breath, irregular breath, after exercise, and sleep. The approach provided a waveform that was 90 percent as accurate as that of a standard ECG.

The researchers describe their work in a study published in October in *IEEE Transactions on Mobile Computing*.

—MICHELLE HAMPSON

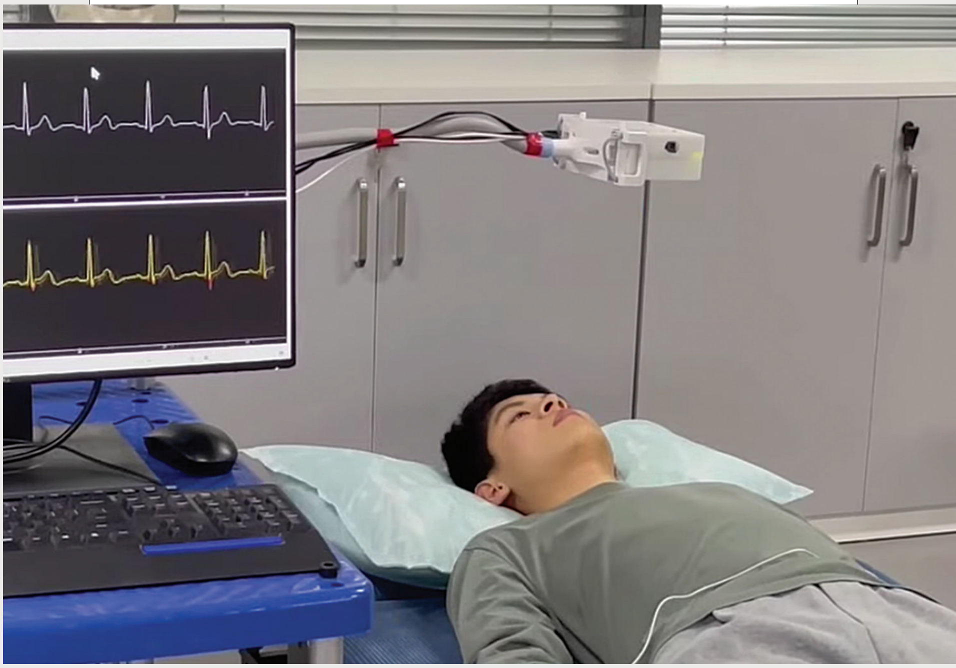

UNIVERSITY OF SCIENCE AND TECHNOLOGY OF CHINA/IEEE

ENERGY

Harnessing Solar Power's Unused Spectrum >

Nanocrystals could convert otherwise wasted infrared light into usable energy

BY PRACHI PATEL

IMAGE SOURCE/GETTY IMAGES

A little more than half of the sun's energy reaches Earth as infrared rays, which today's solar cells can't efficiently capture and convert to electricity. The waste could be avoided if a way were found to turn many low-energy infrared photons into fewer high-energy photons of visible light, a process called upconversion.

Researchers have taken a step in that direction by making semiconductor nanocrystals that can convert infrared light to visible light with an efficiency of more than 5 percent, much higher than methods reported so far. The team plans to use the converted light to drive a chemical reaction that produces hydrogen fuel from methanol.

The research, led by Masanori Sakamoto, a professor of chemistry at Kyoto University, is described in the 17 October 2022 issue of *Nature Sustainability*.

There are two main ways to achieve upconversion. You can shine infrared light onto organic molecules to create unpaired electrons, which then come together to emit higher-energy photons. Or you can shine it onto a metal surface to produce plasmons, a kind of electron wave, which can also recombine to form light.

In the first method, the practical choice of organic materials for solar cells is the lanthanides, such as ytterbium or erbium. When low-energy infrared photons are absorbed by these materials, electrons are excited to ever-higher energy levels until, at last, they fall, emitting high-energy photons. Lanthanides have an upconversion efficiency of 0.5 to 2 percent.

The second method was proposed in 2015 by Gururaj Naik, an electrical and computer engineering professor at Rice University and Jennifer Dionne, a materials scientist at Stanford University. It relies on surface plasmons, waves of electrons that form when light hits a metal surface, appearing as high-energy "hot electrons" and "hot holes."

Naik found a way to channel those hot electrons and holes into structures made of thin alternating layers of indium gallium nitride and gallium nitride. These structures act as quantum wells that trap the hot carriers so they can recombine with holes to release a high-energy photon. But the efficiency was only a puny 0.1 percent, far less than the theoretical maximum of 25 percent.

This is where the Japanese researchers' work comes in. By cleverly designing nanocrystals of cadmium sulfide and copper sulfide, they made quantum wells into which the hot carriers flow more easily and get trapped better. The result is an infrared-to-visible upconversion efficiency of more than 5 percent.

Their nanocrystals are hexagonal, platelike crystals of copper sulfide, with a border of cadmium sulfide. Infrared excitation in the copper sulfide creates the hot electrons and holes, which are injected into the cadmium sulfide and recombine to emit light in the visible region.

Naik, who was not involved in the latest research, says the Japanese team chose the right material and nanoparticle size. "As you shrink the particles, hot-carrier injection gets very efficient," Naik says. "The metal-semiconductor interface also makes a very big difference. These researchers did both—created a better interface and smaller size."

As a demonstration, the researchers placed the nanoparticles in a methanol-water solution and shined infrared light on it. Without the nanoparticles, the low-energy light would do nothing. But the nanoparticles soaked up the infrared light and produced higher-energy light, which was absorbed by platinum catalysts to trigger a chemical reaction that produces hydrogen.

Sakamoto admits that it can be tricky to apply this system to solar cells. So for now, the team is planning to use the innovation to perform other IR-driven photocatalytic reactions that will yield hydrogen and other fuels. ■

The Wurst Use of AI

By Willie D. Jones

From the time the ancient Sumerians started making sausage around 4,000 years ago, the process has been the province of artisans dedicated to the craft of preserving meat so it remained safe to eat for as long as possible. Yet even traditional methods can stand to be improved on from time to time. Katharina Koch of the Landfleischerei Koch in Calden, Germany [right], has retained ancient customs such as curing the Ahle sausages in clay chambers, while also fine-tuning the conditions under which the meats are aged (such as temperature and moisture level) via AI algorithms. The digital modifications she and scientists at the nearby University of Kassel have developed replicate the production methods that have been passed down for generations. So, instead of spending nearly a year manually monitoring the meats' maturation process, a sausage maker using the new AI methods will be able to set it and forget it.

PHOTOGRAPH BY UWE ZUCCHI/ PICTURE ALLIANCE/GETTY IMAGES

Hands On

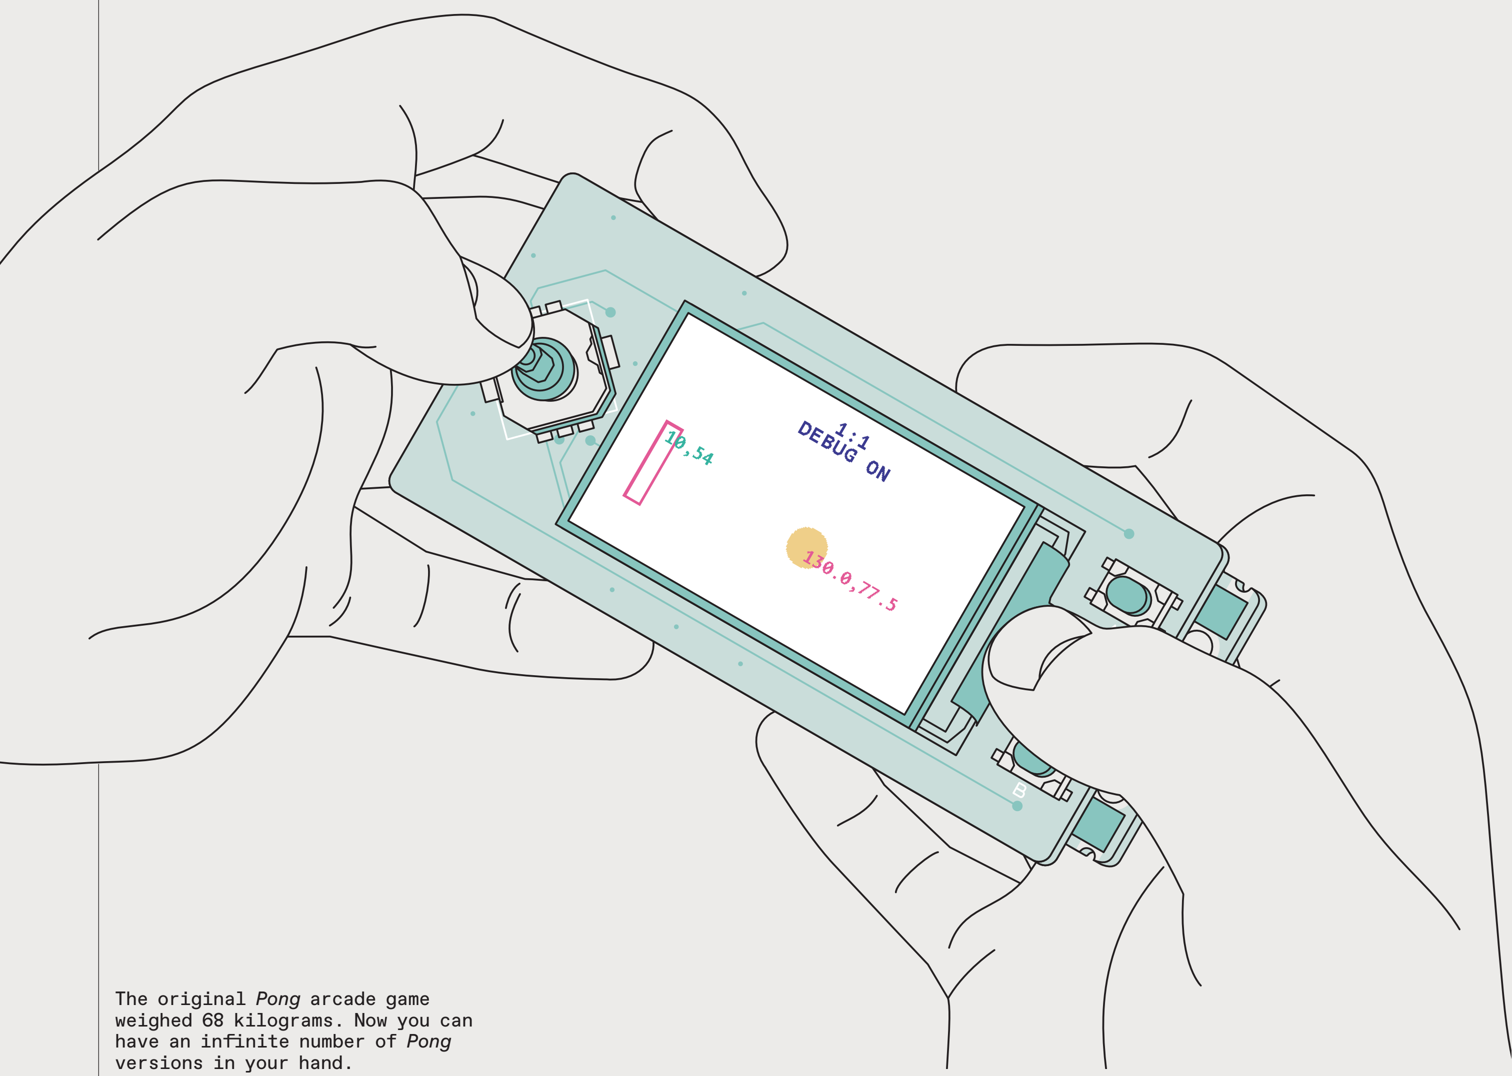

The original *Pong* arcade game weighed 68 kilograms. Now you can have an infinite number of *Pong* versions in your hand.

An Infinity of *Pong* ›

A Raspberry Pi Pico W handheld writes its own games

BY JOSE ANTONIO GARCIA PEIRO

There is currently a lot of interest in AI tools designed to help programmers write software. GitHub's Copilot and Amazon's CodeWhisperer apply deep-learning techniques originally developed for generating natural-language text by adapting it to generate source code. The idea is that programmers can use these tools as a kind of auto-complete on steroids, using prompts to produce chunks of code that developers can integrate into their software.

Looking at these tools, I wondered: Could we take the next step and take the human programmer out

Illustrations by James Provost

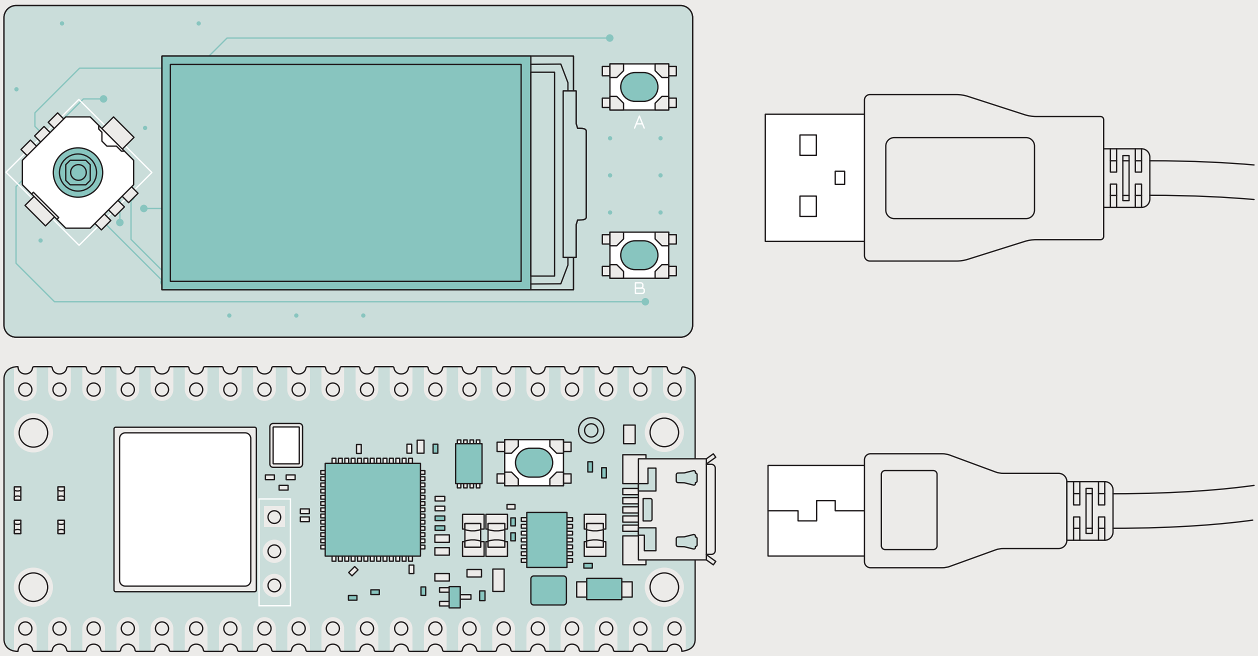

Only two hardware modules are needed—a Raspberry Pi Pico W [bottom left] that supplies the compute power and a plug-in board with a screen and simple controls [top left]. Nothing else is needed except a USB cable to supply power.

of the loop? Could a working program be written and deployed on demand with just the touch of a button?

In my day job, I write embedded software for microcontrollers, so I immediately thought of a self-contained handheld device as a demo platform. A screen and a few controls would allow the user to request and interact with simple AI-generated software. And so was born the idea of infinite *Pong*.

I chose *Pong* for a number of reasons. The gameplay is simple, famously explained on Atari's original 1972 *Pong* arcade cabinet in a triumph of succinctness: "Avoid missing ball for high score." An up button and a down button is all that's needed to play. As with many classic Atari games created in the 1970s and 1980s, *Pong* can be written in a relatively few lines of code, and has been implemented as a programming exercise many, many times. This means that the source-code repositories ingested as training data for the AI tools are rich in *Pong* examples, increasing the likelihood of getting viable results.

I used a US $6 Raspberry Pi Pico W as the core of my handheld device—its built-in wireless allows direct connectivity to cloud-based AI tools. To this I mounted a $9 Pico LCD 1.14 display module. Its 240 x 135 color pixels are ample for *Pong*, and the module integrates two buttons and a two-axis micro joystick.

My choice of programming language for the Pico was MicroPython, because it is what I normally use and because it is an interpreted-language code that can be run without the need of a PC-based compiler. The AI coding tool I used was the OpenAI Codex. The OpenAI Codex can be accessed via an API that responds to queries using the Web's HTTP format. The queries are straightforward to construct and send using the urequests and ujson libraries available for MicroPython. Using the OpenAI Codex API is free during the current beta period, but registration is required and queries are limited to 20 per minute—still more than enough to accommodate even the most fanatical *Pong* jockey.

The next step was to create a container program. This program is responsible for detecting when a new version of *Pong* is requested via a button push, sends a prompt to the OpenAI Codex, receives the results, and launches the game. The container program also sets up a hardware abstraction layer, which handles the physical connection between the Pico and the LCD/control module.

The most critical element of the whole project was creating the prompt that is transmitted to the OpenAI Codex every time we want it to spit out a new version of *Pong*. The prompt is a chunk of plaintext with the barest skeleton of source code—a few lines outlining a CONTINUED ON P. 47

Careers

Technologist Tynesia Boyea-Robinson helps companies incorporate equitable practices into their daily operations.

Tynesia Boyea-Robinson

› She uses a systems-engineering approach to overcoming systemic racism

BY ROBB MANDELBAUM

In parts of the United States, using the term "systemic racism" to refer to persistent discrimination against Black people has become a political flash point. To some ears, it sounds like an attack on the country and the local community. Several states have enacted laws that ban, or would appear to ban, discussing the concept in public schools and colleges, and even private workplaces. But racial-equity consultant Tynesia Boyea-Robinson uses the term with an engineer's precision. When she first heard the phrase, she recalled her training in quality control in the transportation unit of GE Research, in Erie, Pa. And, sure enough, a lightbulb went on in her head: The system could be reengineered. "Oh my God, we can fix this!" she thought. "I don't think everybody else sees it that way."

Boyea-Robinson helps companies, government agencies, and other organizations meet goals for diversity and equity through her consulting firm, CapEQ. In October, her second book on this work, *The Social Impact Advantage*, was published. And she is the steward of Path to 15/55, an ambitious effort to deliver desperately needed capital to Black businesses across the United States. Since 2018, Boyea-Robinson has been assembling a coalition—including financial institutions, grassroots community groups, political and policy leaders, and corporate and philanthropic donors—to reprogram the systems of lending to and investing in these businesses.

Boyea-Robinson grew up in Cocoa Beach, Fla., where her father fixed satellites for the U.S. Air Force and her stepmother gave manicures in the family's living room. In other circumstances, the straight As Boyea-Robinson earned at school and the lessons in mechanics her dad taught her might have ensured a trajectory toward a top STEM university. But her parents hadn't gone to college and didn't push her in that direction. Moreover, as the oldest, she was expected to help care for her four younger siblings. She expected to enroll at a community college until one of her stepmother's clients pushed her to set her sights higher.

She attended Duke University's Pratt School of Engineering, in Durham, N.C., where she earned a dual bachelor's degree in electrical engineering and computer science. The curriculum was daunting, and she had to confront a persistent sense of being an outsider. But it was more than just the academics.

"There's so many things about the culture of college that my parents couldn't teach me," she says. Adding to her initial anxiety was her status as one of the relatively few women at the engineering school—women made up just a quarter of the student body at Pratt—and there were even fewer Black students enrolled there (around 5 percent).

But when Boyea-Robinson graduated in 1999, she landed a plum information-management job at General Electric through the company's prestigious leadership program. Though her anxiety about fitting in lingered, her career flourished. In 2003, she headed to Harvard Business School for an MBA that could give her upward trajectory an

BRANDPRENUER

extra boost. Then her course changed when she took an internship at a nonprofit called Year Up. The organization helps prepare young adults, mostly poorer people of color, for entry-level IT jobs at large companies—jobs that recalled her first assignments at GE. "That student was me," she says, "with different options and choices."

Her assignment was to map out an expansion of Year Up from Boston to either Washington, D.C., or New York City. Boyea-Robinson pitched both. When she graduated in 2005, the nonprofit hired her to open the Washington location. She launched the first class in January 2006, and as she built Year Up's presence in Washington, Boyea-Robinson's work became a model for the organization nationwide, starting in New York later that year. Today, the nonprofit serves 16 metro areas and operates virtually in five others.

At Year Up, Boyea-Robinson began to hear about systemic racism, the biases that people collectively inject, consciously or not, into so many of the institutions and the rules governing society, leading to the disparate treatment of different groups of people. The knock-on effects from that discrimination exacerbate inequality—which then reinforces those biases in a sort of feedback loop. Thinking about all this, Boyea-Robinson concluded that she wanted to use systems engineering to tackle the problems of systemic racism on a larger scale.

Since launching CapEQ in 2011, Boyea-Robinson has worked with more than 50 clients, helping businesses such as Marriott and Nordstrom address their diversity and equity shortcomings. She has also worked with nonprofits and others seeking broader change, including those collaborating on Path to 15/55.

Path to 15/55 takes its premise from recent research by one of those organizations, the Association for Enterprise Opportunity, a trade group of nonprofits that make small loans to underserved entrepreneurs. The group found that if 15 percent of existing Black businesses could finance a single new employee, it would create US $55 billion in new economic activity. But Black entrepreneurs have been hobbled by the effects of an especially pernicious example of systemic racism. Until the 1960s, federal government policies explicitly prohibited Black people from buying homes in white neighborhoods and simultaneously decimated the value of Black neighborhoods. The result has been to deny most Black families the opportunity to build generational wealth on par with their white counterparts. Even today, Blacks are less likely to seek, or obtain, a home mortgage. Most small businesses are financed by savings or loans conditioned on good credit scores and a home that serves as collateral.

The coalition Boyea-Robinson assembled is pressing for systemic change on several levels. It's pushing bankers and the financial industry at large to confront their own biases in lending. It also disseminates novel strategies for financing Black businesses to avoid the barriers that Black borrowers face, such as the use of credit scores to assess creditworthiness. The group will then rigorously collect data on which strategies work and which don't to propagate what's successful. Separately, it's agitating for government policy changes to allow these new strategies to flourish.

Boyea-Robinson manages Path to 15/55 as if she were testing software with a feedback loop of its own. It starts with building awareness around a specific issue and forging alliances, or alignments, with like-minded organizations, which then go to work as communities of action to implement change.

"Everything we learn from communities of action becomes the information that we raise awareness on," she says. "And the loop starts again: awareness, alignment, action. These are all unit tests that become systems tests."

Boyea-Robinson still finds resistance to financing equity among bank loan officers. "The way racism shows up in lending is bankers saying that this work is not investable," she says. "Shifting the narrative is why we spend so much time sharing reports and stories."

Backed with a $250,000 grant from the Walmart Foundation, Path to 15/55 launched its first Community of Action in January. Piggybacking on work led by the Beneficial State Foundation, Boyea-Robinson has recruited five financial institutions to experiment with innovative ways to underwrite loans, and to build durable support within their organizations for the work—which, Boyea-Robinson says, is the only way these changes will stick. These institutions are expected to begin lending money by midyear. To lessen the risk of losses, Path to 15/55 will make the $1 million it has raised so far available for these loans.

And she's joining forces with business accelerators to launch a second community of action, aimed at helping Black entrepreneurs buy existing businesses in corporate supply chains, later this year.

"Being able to kind of turbocharge work that is already compelling," she says, "has been pretty exciting." ■

The Association for Enterprise Opportunity found that if 15 percent of existing Black businesses could finance a single new employee, it would create US $55 billion in new economic activity.

THIS IMPLANT TURNS BRAIN WAVES

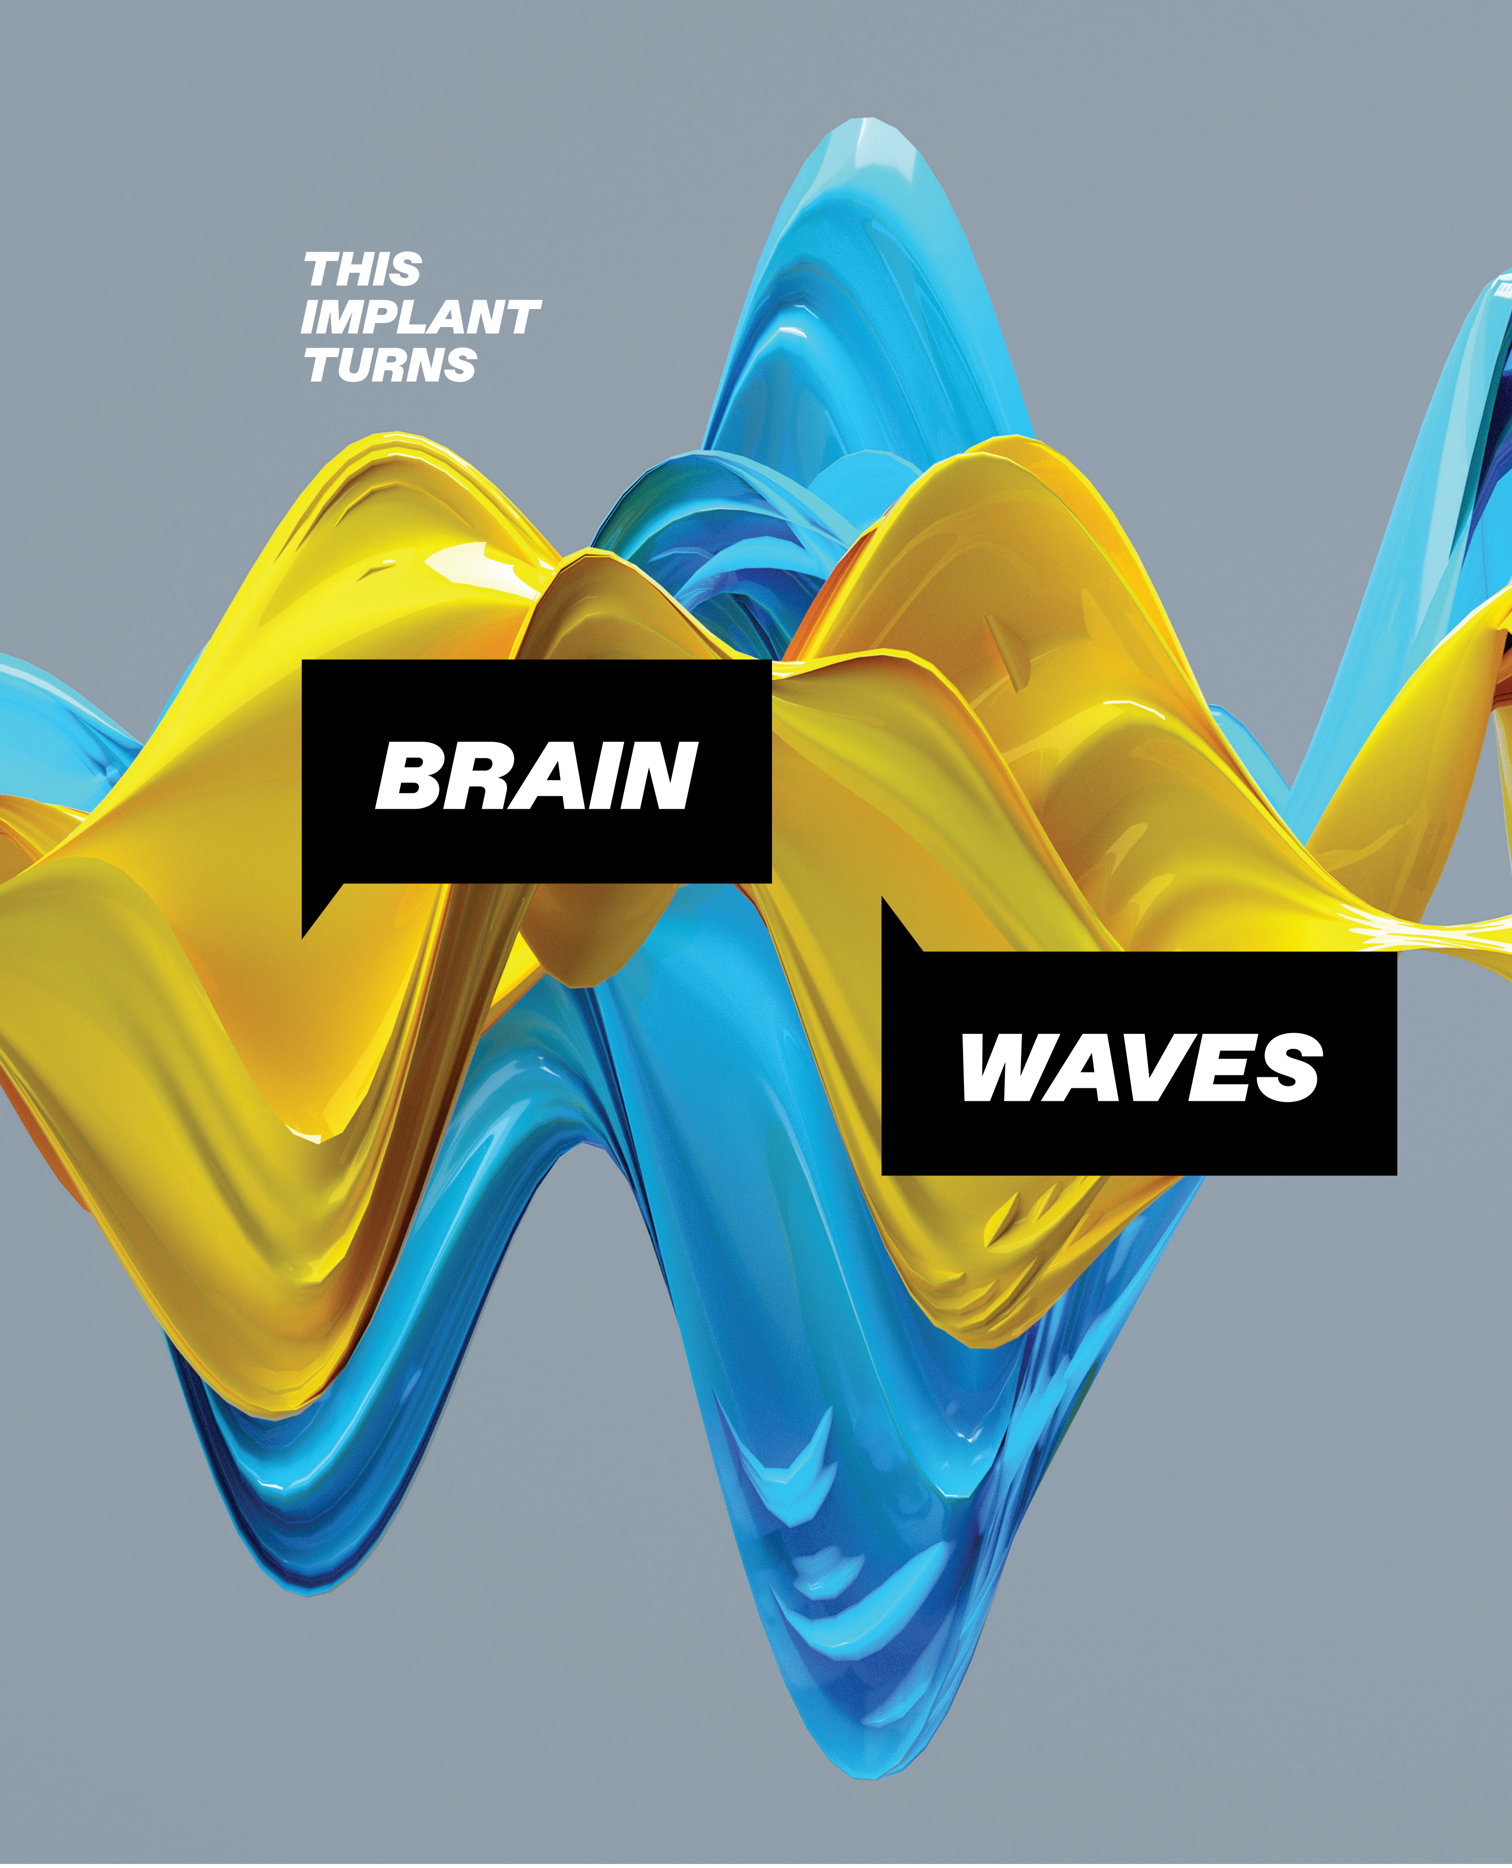

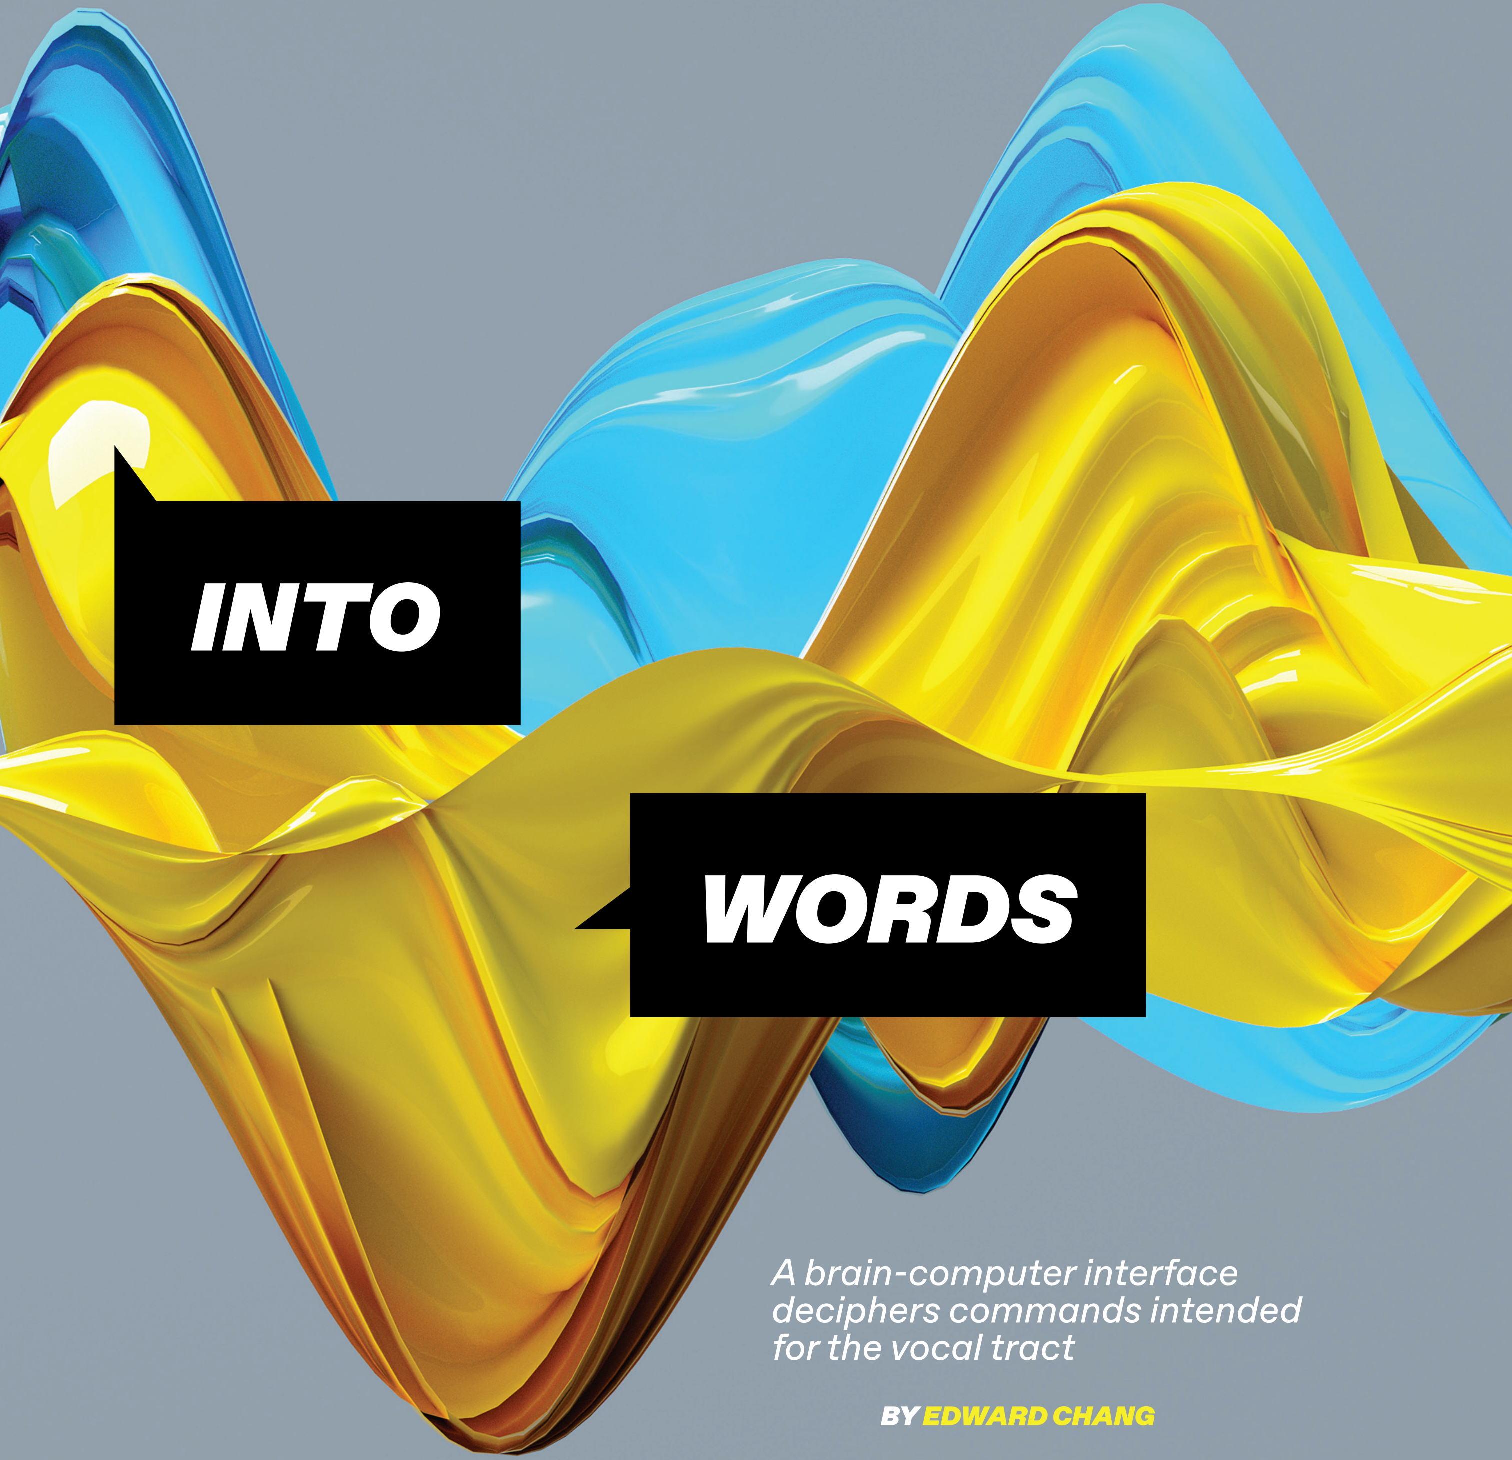

INTO WORDS

A brain-computer interface deciphers commands intended for the vocal tract

BY EDWARD CHANG

A COMPUTER SCREEN SHOWS the question "Would you like some water?" Underneath, three dots blink, followed by words that appear, one at a time: "No I am not thirsty." • It was brain activity that made those words materialize—the brain of a man who has not spoken for more than 15 years, ever since a stroke damaged the connection between his brain and the rest of his body, leaving him mostly paralyzed. He has used many other technologies to communicate; most recently, he used a pointer attached to his baseball cap to tap out words on a touch screen, a method that was effective but slow. He volunteered for my research group's clinical trial at the University of California, San Francisco in hopes of pioneering a faster method. So far, he has used the brain-to-text system only during research sessions, but he wants to help develop the technology into something that people like himself could use in their everyday lives.

In our pilot study, we draped a thin, flexible electrode array over the surface of the volunteer's brain. The electrodes recorded neural signals and sent them to a speech decoder, which translated the signals into the words the man intended to say. It was the first time a paralyzed person who couldn't speak had used neurotechnology to broadcast whole words—not just letters—from the brain.

That trial was the culmination of more than a decade of research on the underlying brain mechanisms that govern speech, and we're enormously proud of what we've accomplished so far. But we're just getting started. My lab at UC San Francisco is working with colleagues around the world to make this technology safe, stable, and reliable enough for everyday use at home. We're also working to improve the system's performance so it will be worth the effort.

NEUROPROSTHETICS HAVE COME a long way in the past two decades. Prosthetic implants for hearing have advanced the furthest, with designs that interface with the cochlear nerve of the inner ear or directly with the auditory brain stem. There's also considerable research on retinal and brain implants for vision, as well as efforts to give people with prosthetic hands a sense of touch. All of these sensory prosthetics take information from the outside world and convert it into electrical signals that feed into the brain's processing centers.

The opposite kind of neuroprosthetic records the electrical activity of the brain and converts it into signals that control something in the outside world, such as a robotic arm, a video-game controller, or a cursor on a computer screen. That last control modality has been used by groups such as the BrainGate consortium to enable paralyzed people to type words—sometimes one letter at a time, sometimes using an auto-complete function to speed up the process.

For that typing-by-brain function, an implant is typically placed in the motor cortex, the part of the brain that controls movement. Then the user imagines certain physical actions to control a cursor that moves over a virtual keyboard. Another approach, pioneered by some of my collaborators in a 2021 paper, had one user imagine that he was holding a pen to paper and writing letters, creating signals in the motor cortex that were translated into text. That approach set a new record for speed, enabling the volunteer to write about 18 words per minute.

In my lab's research, we've taken a more ambitious approach. Instead of decoding a user's intent to move a cursor or a pen, we decode the intent to control the vocal tract, comprising dozens of muscles governing the larynx (commonly called the voice box), the tongue, and the lips.

I began working in this area more than 10 years ago. As a neurosurgeon, I would often see patients with severe injuries that left them unable to speak. To my surprise, in many cases the locations of brain injuries didn't match up with the syndromes I

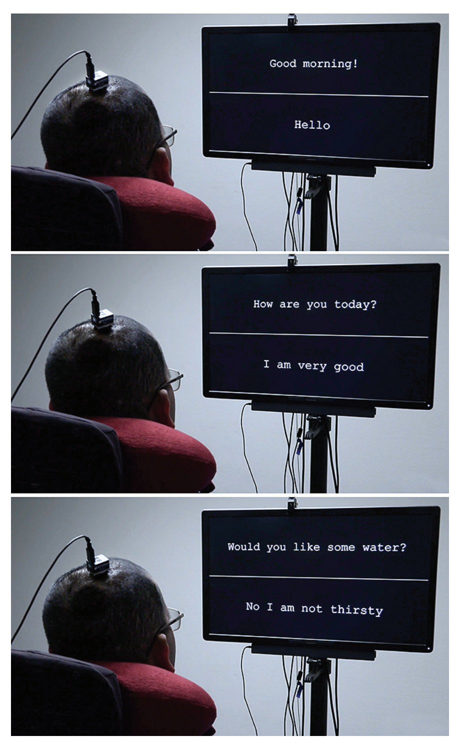

The first version of the brain-computer interface gave the volunteer a vocabulary of 50 practical words.

PREVIOUS PAGES: ANDRIY ONUFRIYENKO/GETTY IMAGES; THIS PAGE: UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

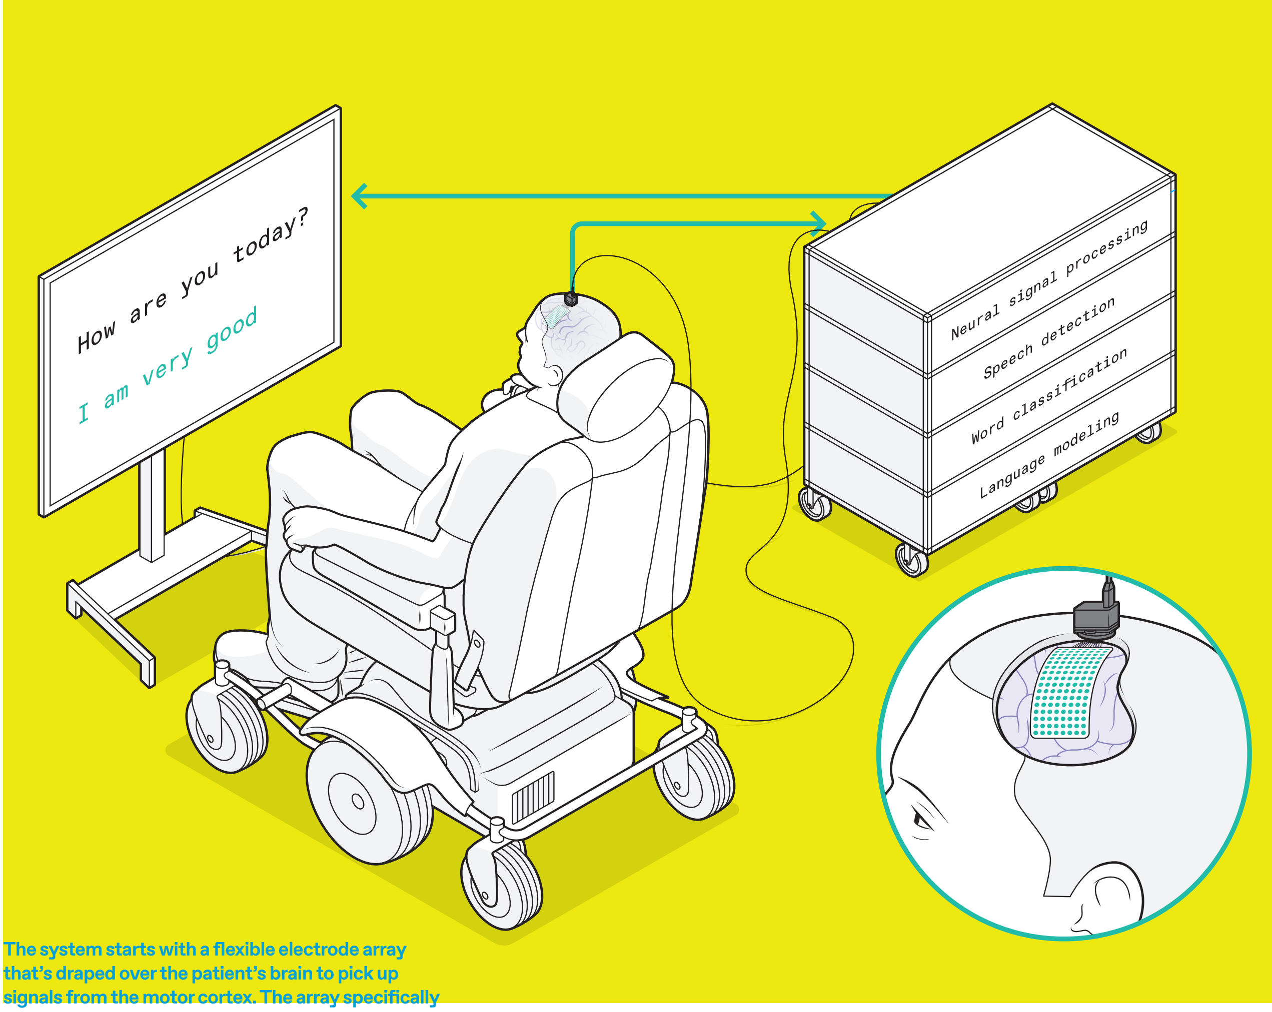

The system starts with a flexible electrode array that's draped over the patient's brain to pick up signals from the motor cortex. The array specifically captures movement commands intended for the patient's vocal tract. A port affixed to the skull guides the wires that go to the computer system, which decodes the brain signals and translates them into the words that the patient wants to say. His answers then appear on the display screen.

learned about in medical school, and I realized that we still have a lot to learn about how language is processed in the brain. I decided to study the underlying neurobiology of language and, if possible, to develop a brain-machine interface (BMI) to restore communication for people who have lost it. In addition to my neurosurgical background, my team has expertise in linguistics, electrical engineering, computer science, bioengineering, and medicine. Our ongoing clinical trial is testing both hardware and software to explore the limits of our BMI and determine what kind of speech we can restore to people.

SPEECH IS ONE of the behaviors that sets humans apart. Plenty of other species vocalize, but only humans combine a set of sounds in myriad different ways to represent the world around them. It's also an extraordinarily complicated motor act—some experts believe it's the most complex motor action that people perform. Speaking is a product of modulated air flow through the vocal tract; with every utterance we shape the breath by creating audible vibrations in our laryngeal vocal folds and changing the shape of the lips, jaw, and tongue.

Many of the muscles of the vocal tract are quite unlike the joint-based muscles in the arms and legs, which can move in only a few prescribed ways. For example, the muscle that controls the lips is a sphincter, while the muscles that make up the tongue are governed more by hydraulics—the tongue is largely composed of a fixed volume of muscular tissue, so moving one part of the tongue changes its shape elsewhere. The physics governing the movements of such muscles is totally different from that of the biceps or hamstrings.

Because there are so many muscles involved and they each have so many degrees of freedom, there's essentially an infinite number of possible configurations. But when people speak, it turns out they use a relatively small set of core movements (which differ somewhat in different languages). For example, when English speakers make the "d" sound, they put their tongues behind their teeth; when they make the "k" sound, the backs of their tongues go up to touch the roof of the back of the mouth. Few people are conscious of the precise, complex, and coordinated muscle actions required to say the simplest word.

My research group focuses on the parts of the brain's motor cortex that send movement commands to the muscles of the face, throat, mouth, and tongue. Those brain regions are multitaskers: They manage muscle movements that produce speech and also the movements of those same muscles for swallowing, smiling, and kissing.

Illustration by Chris Philpot

Studying the neural activity of those regions in a useful way requires both spatial resolution on the scale of millimeters and temporal resolution on the scale of milliseconds. Historically, noninvasive imaging systems have been able to provide one or the other, but not both. When we started this research, we found remarkably little data on how brain activity patterns were associated with even the simplest components of speech: phonemes and syllables.

Here we owe a debt of gratitude to our volunteers. At the UC San Francisco epilepsy center, patients preparing for surgery typically have electrodes surgically placed over the surfaces of their brains for several days so we can map the regions involved when they have seizures. During those few days of wired-up downtime, many patients volunteer for neurological research experiments that make use of the electrode recordings from their brains. My group asked patients to let us study their patterns of neural activity while they spoke words.

The technique involved is called electrocorticography (ECoG). The electrodes in an ECoG system don't penetrate the brain but lie on the surface of it. Our arrays can contain several hundred electrode sensors, each of which records from thousands of neurons. So far, we've used an array with 256 channels. Our goal in the early studies was to discover the patterns of cortical activity when people speak simple syllables. We asked volunteers to say specific sounds and words while we recorded their neural patterns and tracked the movements of their tongues and mouths. Sometimes we did so by having them wear colored face paint and using a computer-vision system to extract the kinematic gestures; other times we used an ultrasound machine positioned under the patients' jaws to image their moving tongues.

We used these systems to match neural patterns to movements of the vocal tract. At first we had a lot of questions about the neural code. One possibility was that neural activity encoded directions for particular muscles, and the brain essentially turned these muscles on and off as if pressing keys on a keyboard. Another idea was that the code determined the velocity of the muscle contractions. Yet another was that neural activity corresponded with coordinated patterns of muscle contractions used to produce a certain sound. (For example, to make the "aaah" sound, both the tongue and the jaw need to drop.) What we discovered was that there is a map of representations that controls different parts of the vocal tract, and that together the different brain areas combine in a coordinated manner to give rise to fluent speech.

OUR WORK DEPENDS on the advances in artificial intelligence over the past decade. We can feed the data we collected about both neural activity and the kinematics of speech into a neural network, then let the machine-learning algorithm find patterns in the associations between the two data sets. It was possible to make connections between neural activity and produced speech, and to use this model to produce computer-generated speech or text. But this technique couldn't train an algorithm for paralyzed people because we'd lack half of the data: We'd have

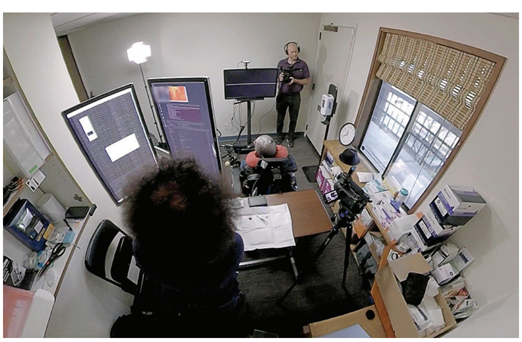

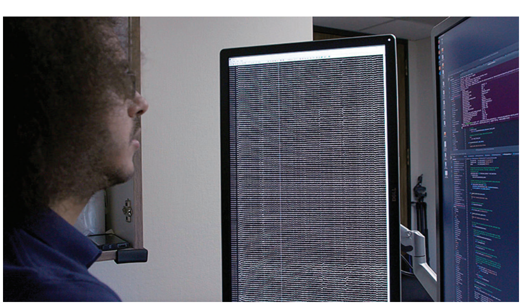

The seemingly simple conversational setup for the paralyzed man [top left, center of photo, in pink shirt] is enabled by sophisticated neurotech hardware and machine-learning systems that decode his brain signals. The author, Edward Chang [above], was inspired to develop the system by patients in his neurosurgery practice. Team member David Moses [bottom left] looks at a readout of the patient's brain waves [left screen] and a display of the decoding system's activity [right screen].

CLOCKWISE FROM RIGHT: BARBARA RIES; UNIVERSITY OF CALIFORNIA, SAN FRANCISCO (2)

Research has shown that speed of setup and performance reliability are key to getting people to use neural devices.

the neural patterns, but nothing about the corresponding muscle movements.

The smarter way to use machine learning, we realized, was to break the problem into two steps. First, the decoder translates signals from the brain into intended movements of muscles in the vocal tract, then it translates those intended movements into synthesized speech or text.

We call this a biomimetic approach because it copies biology; in the human body, neural activity is directly responsible for the vocal tract's movements and is only indirectly responsible for the sounds produced. A big advantage of this approach comes in the training of the decoder for that second step of translating muscle movements into sounds. Because those relationships between vocal tract movements and sound are fairly universal, we were able to train the decoder on large data sets derived from people who weren't paralyzed.

THE NEXT BIG CHALLENGE was to bring the technology to the people who could really benefit from it.

The National Institutes of Health (NIH) is funding our pilot trial, which began in 2021. We already have two paralyzed volunteers with implanted ECoG arrays, and we hope to enroll more in the coming years. The primary goal is to improve their communication, and we're measuring performance in terms of words per minute. An average adult typing on a full keyboard can type 40 words per minute, with the fastest typists reaching speeds of more than 80 words per minute.

We think that tapping into the speech system can provide even better results. Human speech is much faster than typing: An English speaker can easily say 150 words in a minute. We'd like to enable paralyzed people to communicate at a rate of 100 words per minute. We have a lot of work to do to reach that goal, but we think our approach makes it a feasible target.

The implant procedure is routine. First the surgeon removes a small portion of the skull; next, the flexible ECoG array is gently placed across the surface of the cortex. Then a small port is fixed to the skull bone and exits through a separate opening in the scalp. We currently need that port, which attaches to external wires to transmit data from the electrodes, but we hope to make the system wireless in the future.

We've considered using penetrating microelectrodes, because they can record from smaller neural populations and may therefore provide more detail about neural activity. But the current hardware isn't as robust and safe as ECoG for clinical applications, especially over many years.

Another consideration is that penetrating electrodes typically require daily recalibration to turn the neural signals into clear commands, and research on neural devices has shown that speed of setup and performance reliability are key to getting people to use the technology. That's why we've prioritized stability in creating a "plug and play" system for long-term use. We conducted a study looking at the variability of a volunteer's neural signals over time and found that the decoder performed better if it used data patterns across multiple sessions and multiple days. In machine-learning terms, we say that the decoder's "weights" carried over, creating consolidated neural signals.

Because our paralyzed volunteers can't speak while we watch their brain patterns, we asked our first volunteer to try two different approaches. He started with a list of 50 words that are handy for daily life, such as "hungry," "thirsty," "please," "help," and "computer." During 48 sessions over several months, we sometimes asked him to just imagine saying each of the words on the list, and sometimes asked him to overtly *try* to say them. We found that attempts to speak generated clearer brain signals and were sufficient to train the decoding algorithm. Then the volunteer could use those words from the list to generate sentences of his own choosing, such as "No I am not thirsty."

We're now pushing to expand to a broader vocabulary. To make that work, we need to continue to improve the current algorithms and interfaces, but I am confident those improvements will happen in the coming months and years. Now that the proof of principle has been established, the goal is optimization. We can focus on making our system faster, more accurate, and—most important—safer and more reliable. Things should move quickly now.

Probably the biggest breakthroughs will come if we can get a better understanding of the brain systems we're trying to decode, and how paralysis alters their activity. We've come to realize that the neural patterns of a paralyzed person who can't send commands to the muscles of their vocal tract are very different from those of an epilepsy patient who can. We're attempting an ambitious feat of BMI engineering while there is still lots to learn about the underlying neuroscience. We believe it will all come together to give our patients their voices back. ■

WHO *REALLY* INVENTED THE THUMB DRIVE?

THE UBIQUITOUS USB GADGET IS THE BRAINCHILD OF AN UNSUNG SINGAPORE INVENTOR

By Hallam Stevens | *Photography by* Maurizio Di Iorio

Henn Tan, shown here in 2017, fought a series of mostly losing battles against those who pirated Trek 2000's ThumbDrive design and against rival patent claims.

In 2000, at a trade fair in Germany, an obscure Singapore company called Trek 2000 unveiled a solid-state memory chip encased in plastic and attached to a Universal Serial Bus (USB) connector. The gadget, roughly the size of a pack of chewing gum, held 8 megabytes of data and required no external power source, drawing power directly from a computer when connected. It was called the ThumbDrive. ¶ That device, now known by a variety of names—including memory stick, USB stick, flash drive, as well as thumb drive—changed the way computer files are stored and transferred. Today it is familiar worldwide.

The thumb drive was an instant hit, garnering hundreds of orders for samples within hours. Later that year, Trek went public on the Singapore stock exchange, and in four months—from April through July 2000—it manufactured and sold more than 100,000 ThumbDrives under its own label.

Before the invention of the thumb drive, computer users stored and transported their files using floppy disks. Developed by IBM in the 1960s, first 8-inch and later 5¼-inch and 3½-inch floppy disks replaced cassette tapes as the most practical portable storage media. Floppy disks were limited by their relatively small storage capacity—even double-sided, double-density disks could store only 1.44 MB of data.

During the 1990s, as the size of files and software increased, computer companies searched for alternatives. Personal computers in the late 1980s began incorporating CD-ROM drives, but initially these could read only from prerecorded disks and could not store user-generated data. The Iomega Zip Drive, called a "superfloppy" drive and introduced in 1994, could store up to 750 MB of data and was writable, but it never gained widespread popularity, partly due to competition from cheaper and higher-capacity hard drives.

Computer users badly needed a cheap, high-capacity, reliable, portable storage device. The thumb drive was all that—and more. It was small enough to slip in a front pocket or hang from a key chain, and durable enough to be rattled around in a drawer or tote without damage. With all these advantages, it effectively ended the era of the floppy disk.

But Trek 2000 hardly became a household name. And the inventor of the thumb drive and Trek's CEO, Henn Tan, did not become as famous as other hardware pioneers like Robert Noyce, Douglas Engelbart, or Steve Jobs. Even in his home of Singapore, few people know of Tan or Trek.

Why aren't they more famous? After all, mainstream companies including IBM, TEAC, Toshiba, and, ultimately, Verbatim licensed Trek's technology for their own memory stick devices. And a host of other companies just copied Tan without permission or acknowledgment.

The story of the thumb drive reveals much about innovation in the sili-

THIS PAGE: YEN MENG JIIN/SINGAPORE PRESS/AP; OPPOSITE PAGE: NICKY LOH/REUTERS/ALAMY con age. Seldom can we attribute inventions in digital technology to one individual or company. They stem instead from tightly knit networks of individuals and companies working cooperatively or in competition, with advances made incrementally. And this incremental nature of innovation means that controlling the spread, manufacturing, and further development of new ideas is almost impossible.

So it's not surprising that overlapping and competing claims surround the origin of the thumb drive.

In April 1999, the Israeli company M-Systems filed a patent application titled "Architecture for a Universal Serial Bus-based PC flash disk." This was granted to Amir Ban, Dov Moran, and Oron Ogdan in November 2000. In 2000, IBM began selling M-Systems' 8-MB storage devices in the United States under the less-than-memorable name DiskOnKey. IBM has its own claim to the invention of an aspect of the device, based on a year-2000 confidential internal report written by one of its employees, Shimon Shmueli. Somewhat less credibly, inventors in Malaysia and China have also claimed to be the first to come up with the thumb drive.

The necessary elements were certainly ripe for picking in the late 1990s. Flash memory became cheap and robust enough for consumer use by 1995. The circulation of data via the World Wide Web, including software and music, was exploding, increasing a demand for portable data storage.

When technology pushes and consumers pull, an invention can seem, in retrospect, almost inevitable. And all of the purported inventors could certainly have come up with the same essential device independently. But none of the many independent stories of invention paint quite as clear an origin story—or had as much influence on the spread of the thumb drive—as the tale of Tan in Singapore.

Henn Tan holds up a ThumbDrive during an interview in Singapore in January 2006.

TAN, THE THIRD OF SIX BROTHERS, was born and raised in a *kampung* (village) in the neighborhood of Geylang, Singapore. His parents, working hard to make ends meet, regularly left Tan and his brothers alone to roam the streets.

The first in his family to attend high school, Tan quickly fell in with a rebellious crowd, skipping school to hang out at roadside *sarabat* (drink) stalls, dressed in "shaggy embroidered jeans, imbibing coffee and cigarettes, and tossing his long mane as he polemicized about rock music and human rights," according to a 2001 article in the *Straits Times*. After a caning for truancy in his third year of high school that served as a wake-up call, Tan settled down to his studies and completed his O-level exams. He entered the National Service in 1973 as a military police instructor, and after serving the required two years, he took a job as a machinist at a German multinational firm.

This wasn't a rare job at the time. In the late 1960s Singapore had embarked on a crash program of industrialization, offering incentives to multinational companies, especially in such high-tech fields as electronics and semiconductors, to set up factories on the island. By the early 1970s, Singapore was home to manufacturing plants for Fairchild Semiconductor, General Electric, Hewlett Packard, and Texas Instruments, among others. These companies were joined by the Japanese firms Matsushita (now Panasonic) in 1973 and Nippon Electric Co. (now NEC) in 1977.

Tan diligently saved money to pay for driving lessons. As soon as he had his license, NEC's semiconductors division hired him as a sales executive. Three years later, in 1980, he moved to Sanyo as a regional sales manager. Over the next 15 years, he rose to the rank of sales director, accumulating a wealth of experience in the electronics industry, including connections to a range of suppliers and customers.

IN 1995, TAN RESIGNED FROM SANYO and purchased Trek, a small, family-run electronics component trading firm in his old neighborhood of Geylang, for just shy of US $1 million. He planned to develop products to license or sell to one or more of the many large multinationals in Singapore.

Meanwhile, worldwide sales of computer equipment had started to boom. Although personal computers and various portable computers had been around since the late 1970s, both Apple and IBM released flagship laptops in 1991 and 1992, respectively. Along with the popularity of laptops came a growing demand for peripherals such as displays, modems, printers, keyboards, mice, graphics adapters, hard drives, CD-ROM drives, and floppy drives. The dot-com boom of 1995 to 2000 further increased demand for personal computing gear.

Many of these electronics products, as well as the chips in them, were produced in Asia, including Hong Kong, Indonesia, Malaysia, South Korea, Taiwan, Thailand—and Singapore—under the OEM system. These original equipment manufacturers made computers for Apple, Dell, and other companies who outsourced the production of their designs.

By the mid-1990s, Singapore had become an important hub for electronics manufacturing, including hard drives and semiconductor wafers, and the island had a significant and growing electronics ecosystem with design and production expertise.

ALL THIS ACTIVITY, however, did not create an easy path for Tan. Many of his old contacts from Sanyo wouldn't do business with a no-name like Trek. And few talented engineers wanted to work for a company that seemed to offer little guarantee of long-term employment. But Tan persisted, and after two years, in 1998, he got his big break: Toshiba Electronics in Singapore appointed Trek as an official design house, an agreement through which Trek would design and manufacture products to be sold under the Toshiba label.

In particular, Toshiba wanted an MP3 player, a compact and portable solid-state device that could copy music files from a computer, to which it would be connected via a USB plug, and then play the music back. Though this was before Apple's 2001 iPod made these devices popular worldwide, a number of MP3 players of varied quality were already on the market in the late 1990s.

As the originator of flash memory, Toshiba manufactured storage chips used in personal computers, laptops, and digital cameras. Toshiba also made portable radios and boom boxes. It wasn't odd that the company wanted to jump into the MP3-player fray.

But Tan reasoned that "if the company just manufactured the player, it would not make a lot of money," according to a 2005 article in the *Straits Times*. Tan thought that by leaving out the ability to play music, the device would become more versatile, able to handle not just MP3s but also text, spreadsheets, images—any kind of computer file. Many companies were already selling music players, but a cheap, USB-driven, versatile storage device might have an even bigger market, Tan suspected, and he could be first to tap it.

Tan did give Toshiba its music player. But he also set his engineers to work on a product that was essentially a music player without the player. The result was the thumb drive.

GETTING TO A WORKING PRODUCT was not trivial—the drive required not only the appropriate combination of hardware but also specially designed firmware that allowed the solid-state storage to interact with a variety of computer operating systems.

But the thumb drive, with its flash memory and USB interface, was hardly a completely novel invention. Tan did not invent flash memory, which was the brainchild of Toshiba engineer Fujio Masuoka in 1980. Nor did he invent the USB port, which had been around since 1996. What was novel was the combination of the USB with flash memory plus a controller and appropriate firmware, all sealed into a plastic case to make a marketable consumer product.

Local circumstances can partly explain why the thumb drive came to be invented where and when it did: Tan's experience at NEC and Sanyo, Trek's contract with Toshiba, and the connections Trek's engineers had made during previous internships at other companies in Singapore were all important. Those same factors, however, also made the invention difficult to control. Once the idea of the thumb drive was out there, many electronics firms immediately set to making their own versions. Tan had filed a patent application for his invention in 2000, a month before the German tech fair where Trek had introduced the device, but a pending patent did little to stop copycats.

In addition to claims by M-Systems and IBM, perhaps the most complicated rivalry came from the Chinese company Netac Technology. It also claimed to have invented the flash memory stick. Cheng Xiaohua and Deng Guoshun had previously worked for Trek and had seen some development boards related to flash memory. They returned to Shenzhen, China, and founded Netac in 1999.

Shenzhen at the time was a hotbed of electronics copycatting—DVD players, cellular phones, MP3 players, and numerous other consumer electronics were produced as "shanzhai" goods, outside the bounds of intellectual property laws. Netac's claim to (and production of) its thumb drive fit this pattern of appropriation.

Netac and Trek subsequently even entered into an agreement under which Trek would fund some of Netac's research and development and Trek would gain rights to manufacture and distribute the resulting products outside of China. Despite this collaboration, Netac sought and was granted a patent on the thumb drive within China.

Electronics pirates around the world then went after the thumb drive. Tan fought them hard and sometimes won. Had Trek been a larger company with more resources and more patent experi-

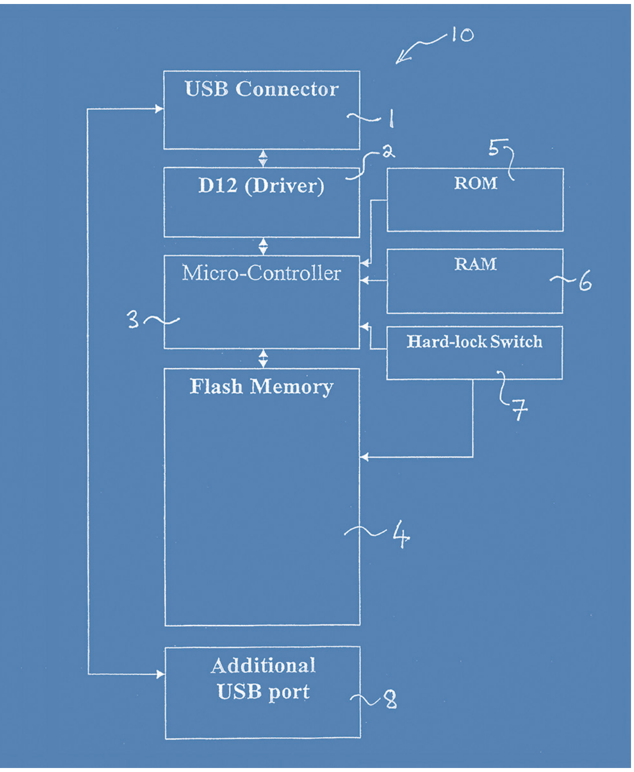

Trek's patent application for the ThumbDrive included this drawing.

INTELLECTUAL PROPERTY OFFICE OF SINGAPORE

ence, the story might have had a different ending. As it was, though, Trek's patents stood on relatively weak ground.

Beginning in 2002, Tan brought suit in Singapore against a handful of companies (including Electec, FE Global Electronics, M-Systems, and Ritronics Components) for patent infringement. After several years of court battles and hundreds of thousands of dollars in legal fees, Trek won that case, persuading the judge that its ThumbDrive was the first device ever designed to be plugged directly into a computer without the need for a cable. An appeals court in the United Kingdom, however, was not persuaded, and Trek lost its patent there in 2008. Tan also pursued, with little success, claims at the United States International Trade Commission against other companies, including Imation, IronKey, Patriot, and Verbatim. But even the decision in Singapore was little more than a moral victory. By the late 2000s, millions of thumb drives had already been produced, by countless companies, without Trek's license.

"Clones," Tan told the *Straits Times* in 2005, "in a sense, are marvelous. In the business world, especially when you are in Asia, as long as anything makes a profit, you do it." If someone were copying you, Tan reasoned, "it meant you must have a good idea and you should make the most of it, as quickly as possible."

Ultimately, Tan and Trek turned their attention to new products, each improving slightly on the last. By 2010, Trek had developed another pioneering device—the Flu Drive or Flu Card. This modified thumb drive could also wirelessly transmit data between devices or to the cloud. Although Tan still attempted to protect his invention with patents, he had also embraced a new path: success through continuous novelty.

The Flu Card enjoyed modest success. Although not widely taken up as a stand-alone device, its Wi-Fi connectivity made it suitable for consumer electronics devices such as cameras and toys. In 2014, Trek signed deals with Ricoh and Mattel China to license the Flu Card design.

Trek also attempted to move into new markets, with limited success, including the Internet of Things, cloud technology, and medical and wearable devices.

TREK'S REVENUE FROM LICENSING the ThumbDrive and the Flu Card was not sufficient to keep it profitable. But instead of admitting how badly the company was doing, in 2006, Tan and his chief financial officer began falsifying Trek's accounts, deceiving auditors and shareholders. After these misdeeds were revealed by financial auditors Ernst & Young in 2015, Tan stepped down as chairman and chief executive and in August 2022 pleaded guilty to falsifying accounts. As of this writing, Tan remains in jail in Singapore. His son, Wayne Tan, continues as Trek's deputy chairman.

Meanwhile, the thumb drive lives on. Although most of us transmit our files over the Internet—either as email attachments or through services like Google Drive and Dropbox—thumb drives (now running to capacities measured in terabytes) remain a convenient device for carrying data in our pockets.

They are used as a quick way to transfer a file from one computer to another, pass out press kits at conferences, lock and unlock computers, carry apps to run on a shared computer, back up travel documents, and even, sometimes, store music. They are used for nefarious purposes as well—to steal files or insert malware into target computers. And they are especially useful for the secure transfer of encrypted data too sensitive to send over the Internet.

In 2021, global sales of the devices from all manufacturers surpassed $7 billion, a number that is expected to rise to more than $10 billion by 2028, according to Vantage Market Research.

OFTEN, WE THINK OF INVENTORS as heroes, boldly going where no one has gone before. But Tan's story isn't that simple.

Tan does deserve a place in consumer electronics history—he conceived the device without seeing one first, made it work, manufactured it in quantities, and spread it broadly, both intentionally through licensing and unintentionally through copying. But full credit for the thumb drive really belongs more to the environment—the ideas circulating at the time and the networks of clients and suppliers—than any individual.

Moreover, the conclusion of Tan's story suggests he is more antihero than hero. We usually admire inventors for their tenacity and grit. In Tan's case, these qualities contributed to his downfall. Determined to take moral and financial credit for the thumb drive, Tan went to extraordinary lengths—even breaking the law—in order to make his company and himself a success. The thumb drive shows how complicated stories of invention often are. ■

The Panopticon v. the Capitol Rioters

Forensic technology is enormously powerful, but is it worth the privacy trade-offs?

by **Mark Harris**

The group of well-dressed young men who gathered on the outskirts of Baltimore on the night of 5 January 2021 hardly looked like extremists. But the next day, prosecutors allege, they would all breach the United States Capitol during the deadly insurrection. Several would loot and destroy media equipment, and one would assault a policeman.

No strangers to protest, the men, members of the America First movement, diligently donned masks to obscure their faces. None boasted of their exploits on social media, and none of their friends or family would come forward to denounce them. But on 5 January, they made one piping hot, family-size mistake: They shared a pizza.

Photo-illustrations by Gabriel Zimmer

According to charging documents, at 10:57 that evening, a PayPal account registered to a Gmail address paid US $84.72 to Domino's Pizza in Arbutus, Md. Minutes later, that email account received Venmo payments from users called Thomas Carey, Gabe Chase, and Jon Lizak. A separate Venmo email showed a payment from "Broseph Broseph," a nickname of another friend, Joseph Brody.

After the horrific events of the next day, the Federal Bureau of Investigation swung into action. It served cell service and tech companies with geofence warrants—search warrants demanding details on every device and app active within a specified geographic area. One of these warrants, served on Google and covering the interior of the Capitol, showed that a device associated with the Gmail account in question entered the Senate Wing door at 2:18 p.m. on 6 January.

Connecting that Gmail account to a phone number and then to its owner, Paul Lovley of Halethorpe, Md., was just a matter of a few keystrokes on law-enforcement databases. All that remained was for an FBI agent on stakeout to observe Lovley taking out the trash one night and match his photo to one of a figure captured by Senate surveillance cameras during the riot. Lovley and his four compatriots were charged with a range of federal crimes in September 2022.

The riot was an unprecedented attack on American democracy, with thousands of citizens, most of them previously unknown to federal investigators, violently storming the seat of government. The resulting investigations were the largest in U.S. history, offering a snapshot of the rapidly evolving nature of law enforcement and how heavily it now relies on data provided, wittingly or not, by suspects themselves.

While it might seem as though the Capitol-riot investigations represent state-of-the-art digital forensics, "those surveillance technologies are being used in even minor low-level criminal cases across the country every single day," says Jennifer Lynch, surveillance litigation director at the Electronic Frontier Foundation (EFF). "The FBI did not use anything new. They just used it at a much larger scale."

IEEE Spectrum analyzed hundreds of criminal complaints and other legal filings from the Capitol attacks to understand that reach and scale, and to consider the legal and social consequences of the government's power to delve into its citizens' digital lives. That power might seem reassuring when applied to a mob intent on overturning a presidential election, but perhaps less so when brought to bear on people protesting, say, human-rights violations.

Social media provides clues for digital forensics

POLICE WORK HAS always involved the connecting of dots, whether photos, phone calls, testimony, or physical evidence. The 6 January investigation showed the power of seeking the digital connections between those dots.

Over the past two years, the U.S. Department of Justice and the Program on Extremism at George Washington University have made available thousands of legal documents about those charged in connection with the 6 January riot. *Spectrum* analyzed all those containing details of how alleged perpetrators were identified and investigated: 884 individuals by mid-December. Many were identified using time-honored techniques: Wanted posters remain a powerful tool, these days reaching a global audience via news organizations, the FBI's website, and social media. Nearly two-thirds of all those people were first identified via tips from witnesses, friends, family, and other human sources. The FBI ultimately received more than 300,000 such tips.

But the ways in which those sources spotted the alleged perpetrators have changed enormously. Only a tiny fraction of sources were on the ground in Washington, D.C., on 6 January. And although some suspects were recognized in TV reports or news stories, most were spotted on social media.

In almost two-thirds of the cases, evidence was cited from one or more social-media platforms. Facebook appeared in almost half

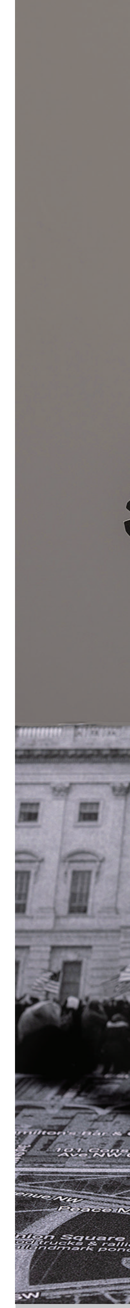

Betrayed by Technology

In nearly 150 of the Capitol riot cases charged, technology provided the initial clue to the suspect. Here's the breakdown:

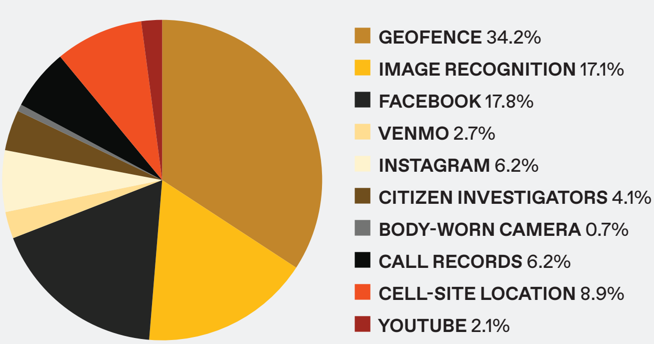

SOURCE: U.S. FEDERAL COURT FILINGS

ORIGINAL PHOTOS FOR ILLUSTRATIONS: GETTY IMAGES

of all cases (388), followed by Instagram and Twitter, which were cited a total of 188 times. But almost every major social-media app was mentioned in at least one case: LinkedIn, MeWe, Parler, Signal, Snapchat, Telegram, TikTok, even dating app Bumble and shopping-focused Pinterest.

Investigators immediately exploited the rioters' use of Facebook. On the day of the attack, the FBI requested that Facebook identify "any users that broadcasted live videos which may have been streamed and/or uploaded to Facebook from physically within the building of the United States Capitol during the time on January 6, 2021, in which the mob had stormed and occupied the Capitol building." Complying with this request was possible because Facebook records the latitude and longitude of every uploaded photo and video by default.

Facebook responded the very same day, and again over the next few weeks, with an unknown number of user IDs—unique identifiers assigned to accounts on Facebook and Instagram (which Facebook's parent company, Meta, also owns). The legal documents suggest that about 35 rioters were identified this way, without first being named by witnesses. In many cases, the FBI then requested that Facebook send it the relevant images and videos and other account data.

Investigators gleaned further clues from many hours of professional news footage, as well as 14,000 hours of high-resolution video from dozens of fixed security cameras and 2,000 hours of video from body-worn cameras operated by police responding to the riot. Surveillance cameras were referenced in 63 percent of DOJ cases, open-source videos and social-media images in 41 percent, and body-camera and news footage each in about 20 percent of cases.

Processing these files involved a huge amount of human effort. The body-camera footage alone required a team of 60, who laboriously completed a 752-page spreadsheet detailing relevant clips.

Shortly after the 6 January riot, *Spectrum* reported on how automated image-recognition systems could be brought to bear on this flood of audiovisual information. The FBI assigned its FACE Services Unit to compare suspects' faces with images in state and federal face-recognition systems. However, according to the legal documents, only 25 rioters appear to have been first identified through such automated image searches, mostly after comparisons with state driver's license photos and passport applications.

Hoan Ton-That, CEO of Clearview AI, a face-recognition search engine that indexes 30 billion images from the open Internet, told *Spectrum* that the court filings do not necessarily reflect how often such technology was used. "Law enforcement don't always have to disclose that they found a certain person's information through facial recognition," he says.

Ton-That notes that Clearview's algorithm is not yet admissible in court, and that any identification it makes from open-source imagery requires further vetting and confirmation. Without providing specifics, he suggested that Clearview's system was used by the FBI. "As a company, it was gratifying for us to play a small role in helping apprehend people who caused damage and stormed the Capitol," he told *Spectrum*. The Capitol riot wouldn't have been the first time that such technology was applied in this way. Facial recognition was reportedly used to identify protesters at a Black Lives Matter event in New York City in 2020 and at similar protests across the United States.

Computers are generally much better at recognizing letters and numbers than faces; automatic license plate reader (ALPR) technology was cited in 20 of the DOJ cases. There are likely tens of thousands of fixed and mobile ALPR systems in the United States alone, at toll plazas, bridge crossings, and elsewhere, capturing hundreds of millions of car journeys each month.

How digital data makes it easier to connect the dots

A SINGLE STREAM of data may help a little, but the integration of many such streams can do wonders. Take the case of William Vogel. He was first named by a tipster who sent the FBI a Snapchat video filmed by someone, unpictured, inside the Capitol building. Sure enough, a Facebook account associated with the Snapchat account listed Vogel as its owner and included a cellphone number.

But maybe someone stole Vogel's cellphone and his Snapchat login to shoot and upload the video. Vogel's phone number led to an address in Pawling, N.Y., and to a car registered to Vogel. The FBI then logged on to ALPR systems across several states, revealing that Vogel's vehicle had taken the Henry Hudson Bridge from the Bronx into Manhattan at 6:06 a.m. on 6 January, entered New Jersey at 7:54, and proceeded southbound through Baltimore at 9:15. The car made its return journey late that afternoon, eventually crossing back into New York a minute before midnight.

But, again, perhaps someone had borrowed Vogel's car? Not according to an ALPR photo snapped in rural Maryland at 8:44 a.m. It shows a distinctive large red "Make America Great Again" hat on the car's dashboard, just like one that Vogel was wearing when he was filmed on a news broadcast outside the Capitol later that day, and in a Facebook selfie.

"They're trying to report me to the FBI/DOJ and put me away for 10 years for domestic terrorism, because of my Snapchat story," Vogel complained later via Facebook Messenger, after admitting to a friend that he had in fact shot the Capitol video, charging documents allege. Vogel's case goes to trial in February 2023, when he will face

Crowds throng the U.S. Capitol Building, in Washington, D.C., on 6 January 2021.

EVELYN HOCKSTEIN/WASHINGTON POST/GETTY IMAGES charges of violently entering the Capitol and disorderly conduct.

Investigators also homed in on people by looking at data from their cellphones. At least 2,000 digital devices were searched by the FBI for images, data, and messages. The FBI's Cellular Analysis Survey Team is dedicated to locating cellphones based on which cell towers they access. Although the FBI got rough locations for about one-fifth of the Capitol-riot defendants this way, it's too imprecise to reliably indicate whether someone actually breached the Capitol itself or remained outside the building.

Far more accurate are the geolocation data gathered by Google Maps and other apps, on both Android and Apple devices. By bolstering cell-tower data with information from nearby Wi-Fi routers and Bluetooth beacons, these apps can locate a target to within about 10 meters (better in urban areas, worse in the countryside). They can even work on phones that have been put in airplane mode.

Until the 6 January attacks, geofence search warrants served on Google—for example, by agents investigating a bank robbery—might produce just a dozen suspect devices. The Capitol breach resulted in 5,723, by far the largest such production. It took until early May 2021 for Google to hand over the data to the FBI; when it did so, the results were comprehensive. That data included the latitude and longitude of each device to seven decimal places, and how long it was inside the Capitol. After narrowing the results to only those most likely to have breached the Capitol, Google eventually delivered the names, phone numbers, and emails associated with the accounts—everything investigators needed to identify and track someone inside the Capitol that day.

And track they did. The legal documents indicate that the Google geofence warrants yielded more initial identifications—50 individuals—than did any other technology, and they were cited in a total of 128 cases. Investigators were able to match interior surveillance footage of one suspect, Raul Jarrin, with a photo he was taking on his Samsung cellphone at the exact same moment. They later acquired the photo from Google under a separate warrant. Jarrin was arrested in March 2022.

On top of the Google data, the FBI served geofence search warrants for anonymized location data from 10 data-aggregation companies. But none of these companies were cited in a criminal complaint, and there are no further details.

The EFF sees the tremendous scope and power of geofence warrants as a bug, not a feature. "We believe that geofence warrants are unconstitutional because they don't start with a suspect," says Lynch. "They don't rely on individualized suspicion, which is what's required under the Fourth Amendment [to the U.S. Constitution]. In the January 6th context, it's likely that there were many journalists whose data was provided to the police."

CONTINUED ON P. 46

Anti–Social Networks

Many accused of crimes on 6 January proudly shared details to their feeds, resulting in tens of thousands of tips to the FBI. Social networks were cited in about two-thirds of investigations. Some cases cited multiple networks.

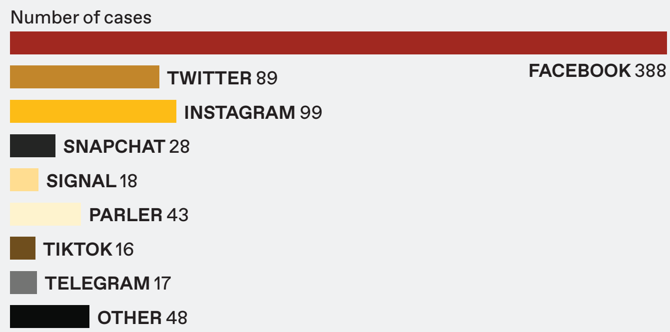

A Picture of Guilt

Video evidence was cited in over 90 percent of the Capitol riot cases. Here's where the pixels originated. Some cases cited multiple video sources.

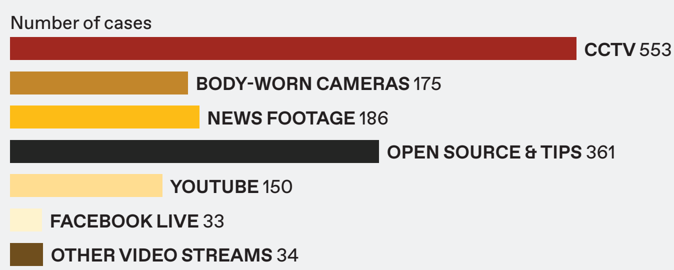

SOURCE: U.S. FEDERAL COURT FILINGS

HIGH-STABILITY FILM RESISTOR

THIS HIGH-STABILITY film resistor, about 4 millimeters in diameter, is made in much the same way as its inexpensive carbon-film cousin, but with exacting precision. A ceramic rod is coated with a fine layer of resistive film (thin metal, metal oxide, or carbon) and then a perfectly uniform helical groove is machined into the film. • Instead of coating the resistor with an epoxy, it's hermetically sealed in a lustrous little glass envelope. This makes the resistor more robust, ideal for specialized cases such as precision reference instrumentation, where long-term stability of the resistor is critical. The glass envelope provides better isolation against moisture and other environmental changes than standard coatings like epoxy.

THE INNER BEAUTY OF BASIC ELECTRONICS

A new book shows the surprising complexity inside passive components

CAPTIONS AND PHOTOS BY
ERIC SCHLAEPFER & WINDELL H. OSKAY

ERIC SCHLAEPFER WAS trying to fix a broken piece of test equipment when he came across the cause of the problem—a troubled tantalum capacitor. The component had somehow shorted out, and he wanted to know why. So he polished it down for a look inside. He never found the source of the short, but he and his collaborator, Windell H. Oskay, discovered something even better: a breathtaking hidden world inside electronics. What followed were hours and hours of polishing, cleaning, and photography that resulted in *Open Circuits: The Inner Beauty of Electronic Components* (No Starch Press, 2022), an excerpt of which follows. As the authors write, everything about these components is deliberately designed to meet specific technical needs, but that design leads to "accidental beauty: the emergent aesthetics of things you were never expected to see." ● From a book that spans the wide world of electronics, what we at *IEEE Spectrum* found surprisingly compelling were the insides of things we don't spend much time thinking about, passive components. Transistors, LEDs, and other semiconductors may be where the action is, but the simple physics of resistors, capacitors, and inductors have their own sort of splendor. —SAMUEL K. MOORE

IT TAKES 15 ROTATIONS of an adjustment screw to move a 15-turn trimmer potentiometer from one end of its resistive range to the other. Circuits that need to be adjusted with fine resolution control use this type of trimmer pot instead of the single-turn variety. • The resistive element in this trimmer is a strip of cermet—a composite of ceramic and metal—silk-screened on a white ceramic substrate. Screen-printed metal links each end of the strip to the connecting wires. It's a flattened, linear version of the horseshoe-shaped resistive element in single-turn trimmers. • Turning the adjustment screw moves a plastic slider along a track. The wiper is a spring finger, a spring-loaded metal contact, attached to the slider. It makes contact between a metal strip and the selected point on the strip of resistive film.

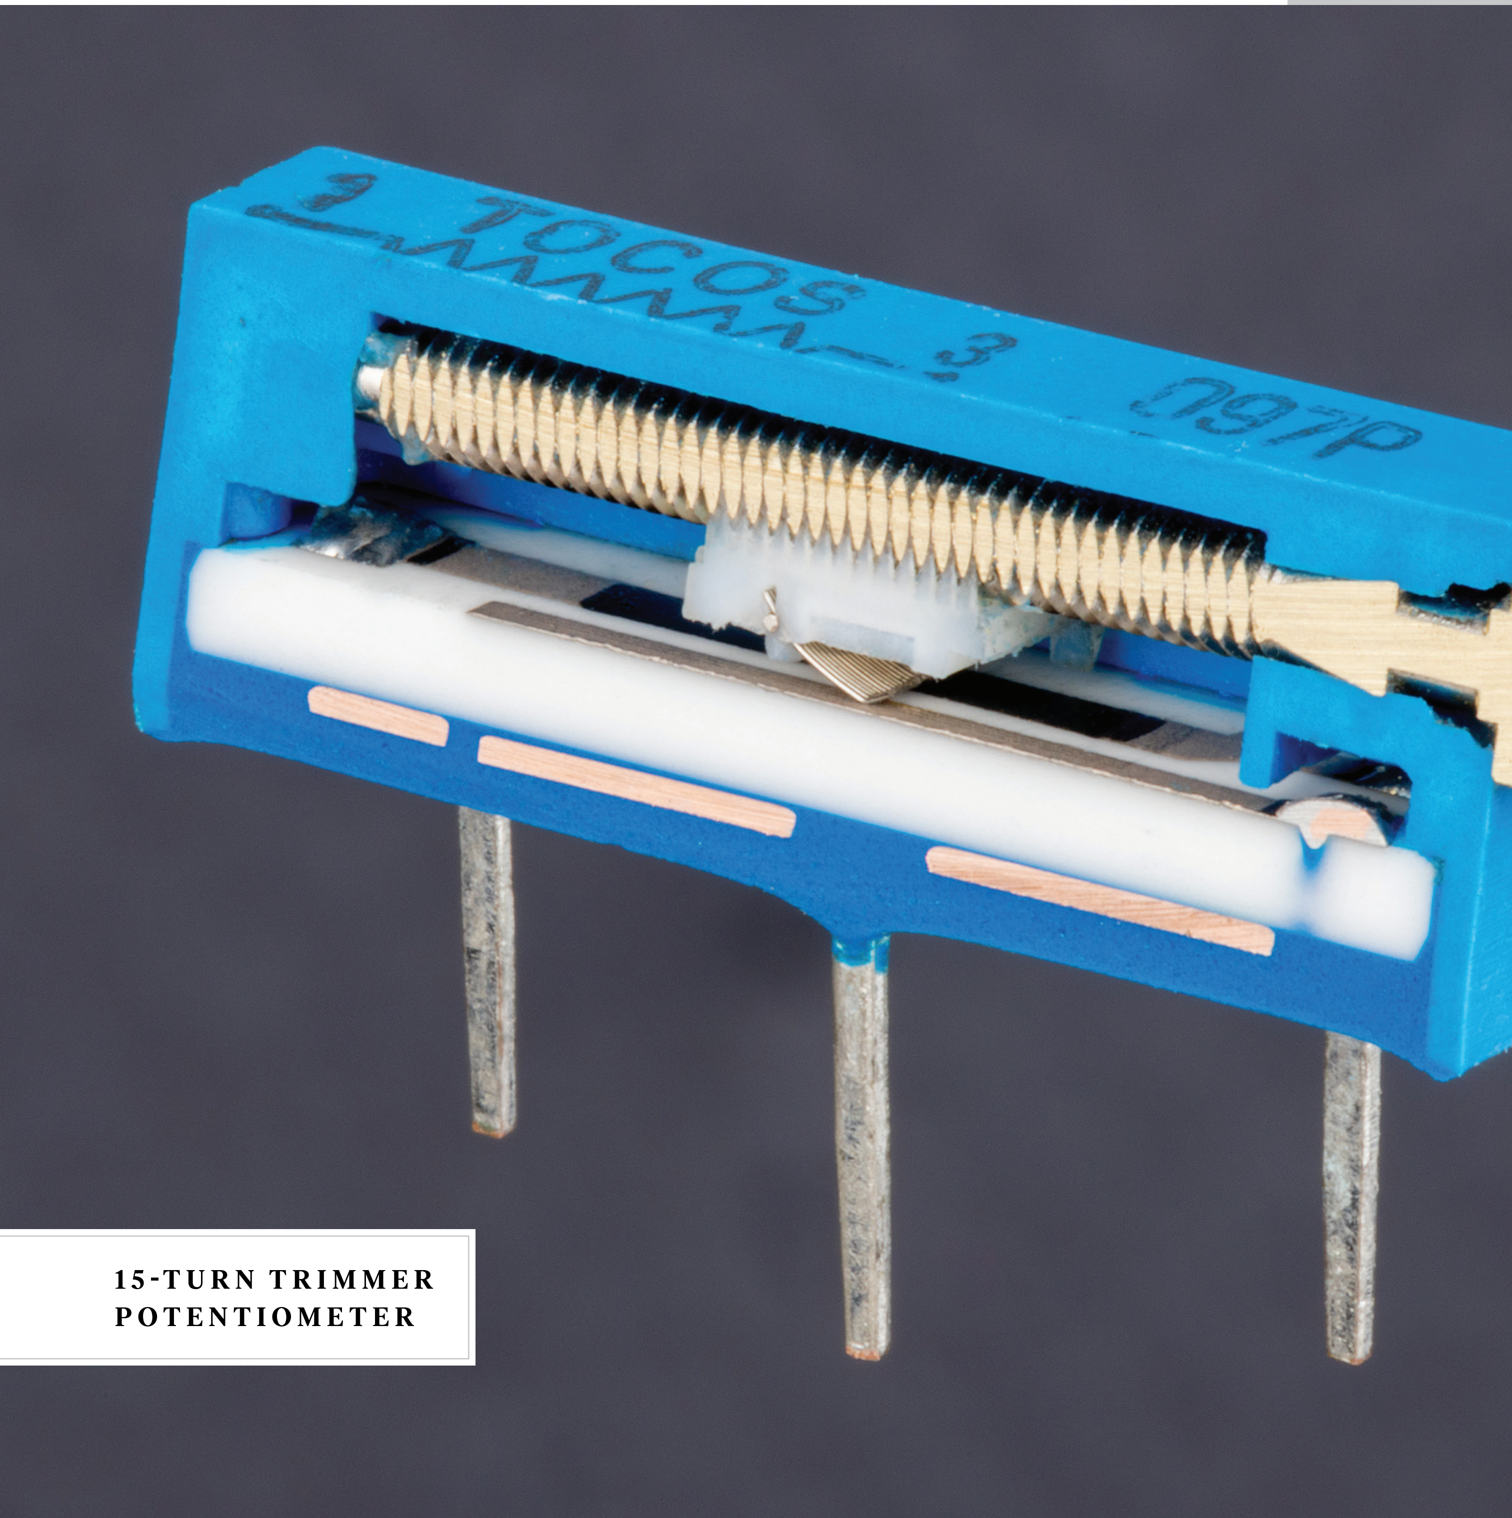

15-TURN TRIMMER POTENTIOMETER

CERAMIC DISC CAPACITOR

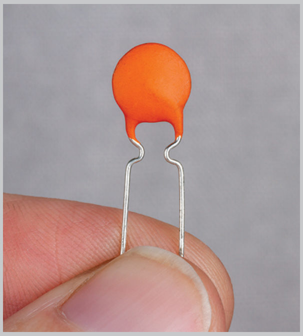

CAPACITORS ARE FUNDAMENTAL electronic components that store energy in the form of static electricity. They're used in countless ways, including for bulk energy storage, to smooth out electronic signals, and as computer memory cells. The simplest capacitor consists of two parallel metal plates with a gap between them, but capacitors can take many forms so long as there are two conductive surfaces, called electrodes, separated by an insulator. • A ceramic disc capacitor is a low-cost capacitor that is frequently found in appliances and toys. Its insulator is a ceramic disc, and its two parallel plates are extremely thin metal coatings that are evaporated or sputtered onto the disc's outer surfaces. Connecting wires are attached using solder, and the whole assembly is dipped into a porous coating material that dries hard and protects the capacitor from damage.

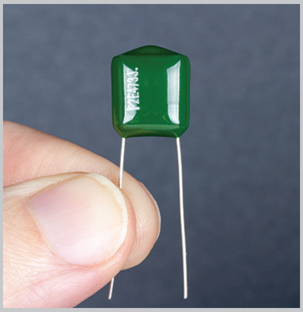

FILM CAPACITORS ARE frequently found in high-quality audio equipment, such as headphone amplifiers, record players, graphic equalizers, and radio tuners. Their key feature is that the dielectric material is a plastic film, such as polyester or polypropylene. • The metal electrodes of this film capacitor are vacuum-deposited on the surfaces of long strips of plastic film. After the leads are attached, the films are rolled up and dipped into an epoxy that binds the assembly together. Then the completed assembly is dipped in a tough outer coating and marked with its value. • Other types of film capacitors are made by stacking flat layers of metallized plastic film, rather than rolling up layers of film.

FILM CAPACITOR

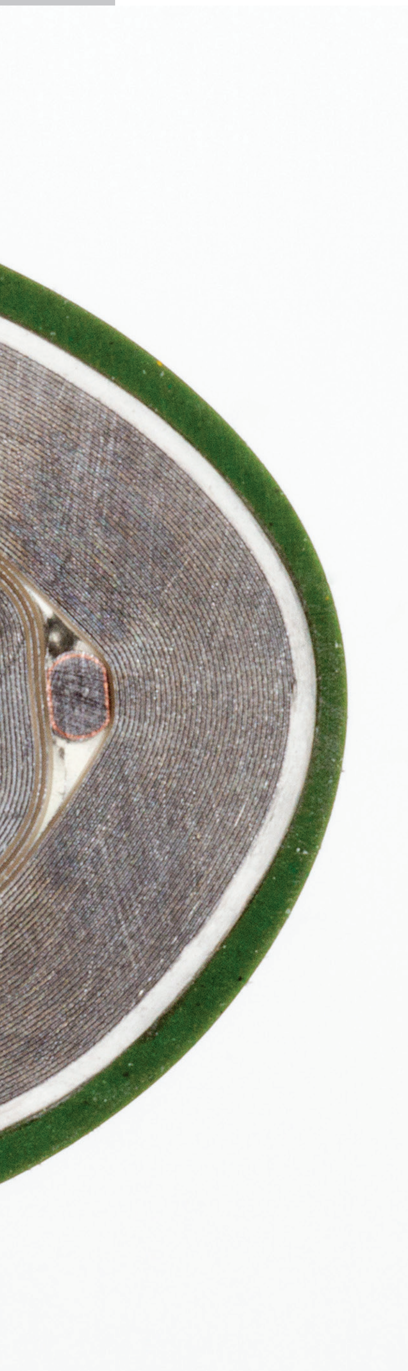

DIPPED TANTALUM CAPACITOR

AT THE CORE of this capacitor is a porous pellet of tantalum metal. The pellet is made from tantalum powder and sintered, or compressed at a high temperature, into a dense, spongelike solid. • Just like a kitchen sponge, the resulting pellet has a high surface area per unit volume. The pellet is then anodized, creating an insulating oxide layer with an equally high surface area. This process packs a lot of capacitance into a compact device, using spongelike geometry rather than the stacked or rolled layers that most other capacitors use. • The device's positive terminal, or anode, is connected directly to the tantalum metal. The negative terminal, or cathode, is formed by a thin layer of conductive manganese dioxide coating the pellet.

POWER SUPPLY TRANSFORMER

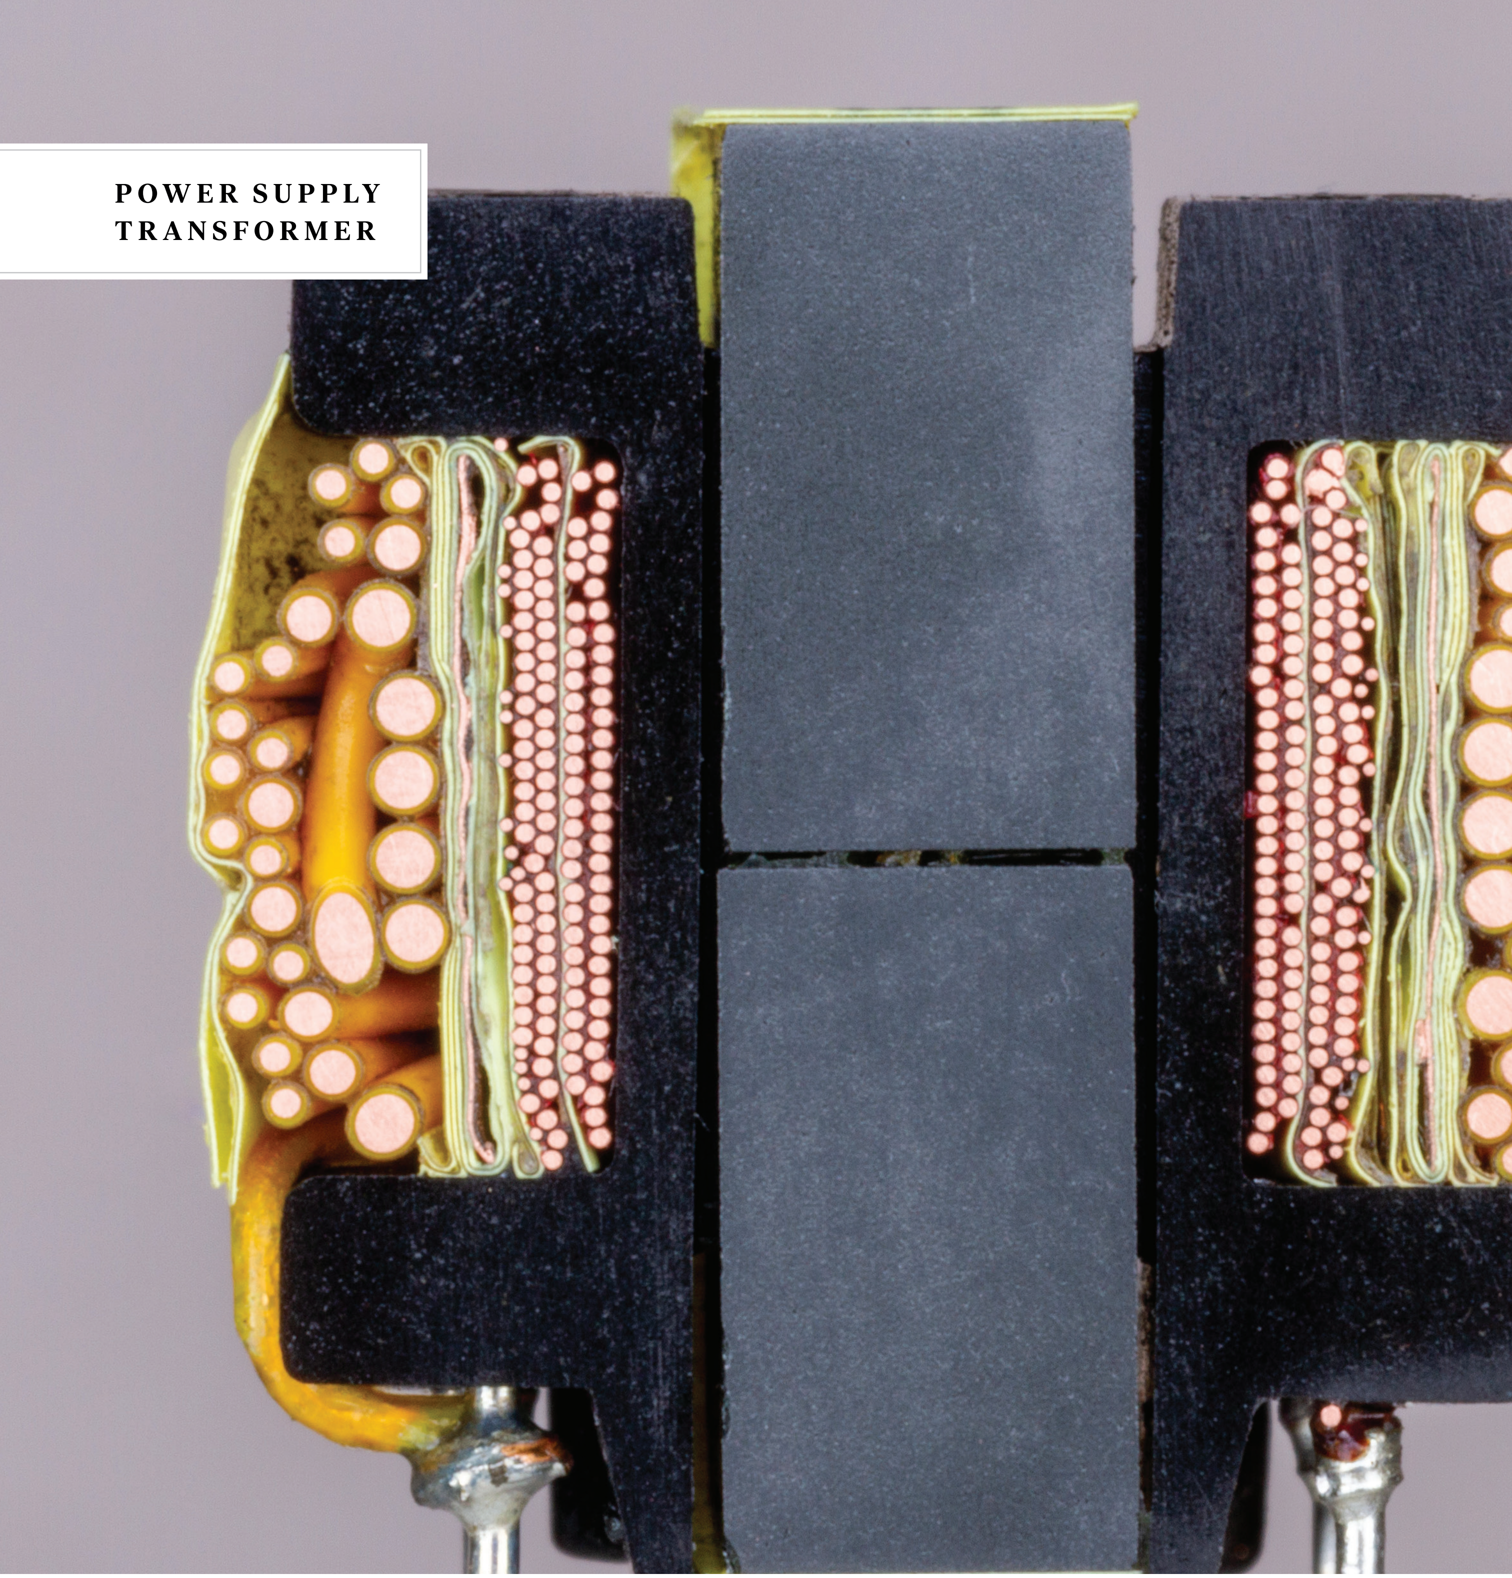

THIS TRANSFORMER HAS multiple sets of windings and is used in a power supply to create multiple output AC voltages from a single AC input such as a wall outlet. • The small wires nearer the center are "high impedance" turns of magnet wire. These windings carry a higher voltage but a lower current. They're protected by several layers of tape, a copper-foil electrostatic shield, and more tape. • The outer "low impedance" windings are made with thicker insulated wire and fewer turns. They handle a lower voltage but a higher current. • All of the windings are wrapped around a black plastic bobbin. Two pieces of ferrite ceramic are bonded together to form the magnetic core at the heart of the transformer.

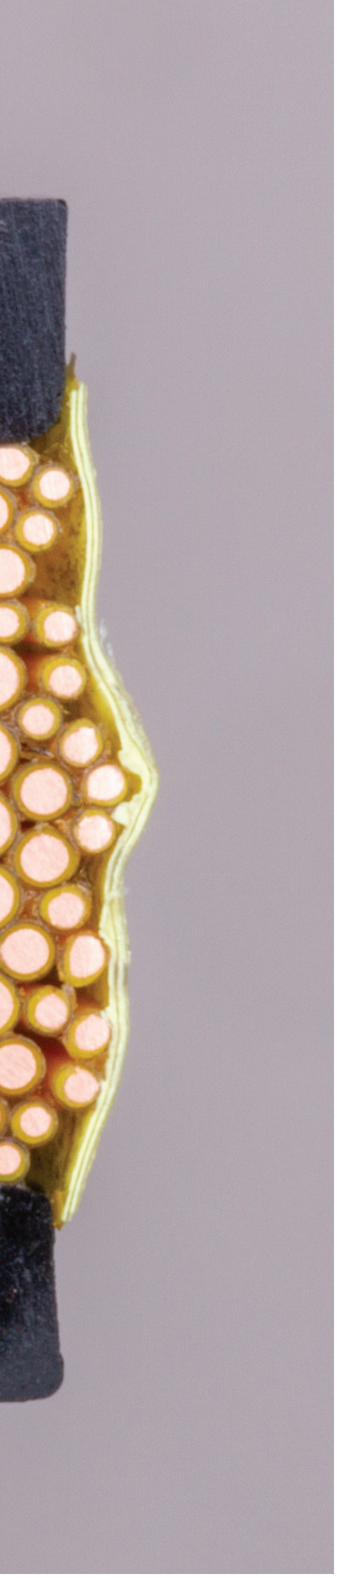

INDUCTORS ARE FUNDAMENTAL electronic components that store energy in the form of a magnetic field. They're used, for example, in some types of power supplies to convert between voltages by alternately storing and releasing energy. This energy-efficient design helps maximize the battery life of cellphones and other portable electronics. • Inductors typically consist of a coil of insulated wire wrapped around a core of magnetic material like iron or ferrite, a ceramic filled with iron oxide. Current flowing around the core produces a magnetic field that acts as a sort of flywheel for current, smoothing out changes in the current as it flows through the inductor. • This axial inductor has a number of turns of varnished copper wire wrapped around a ferrite form and soldered to copper leads on its two ends. It has several layers of protection: a clear varnish over the windings, a light-green coating around the solder joints, and a striking green outer coating to protect the whole component and provide a surface for the colorful stripes that indicate its inductance value.

AXIAL INDUCTOR

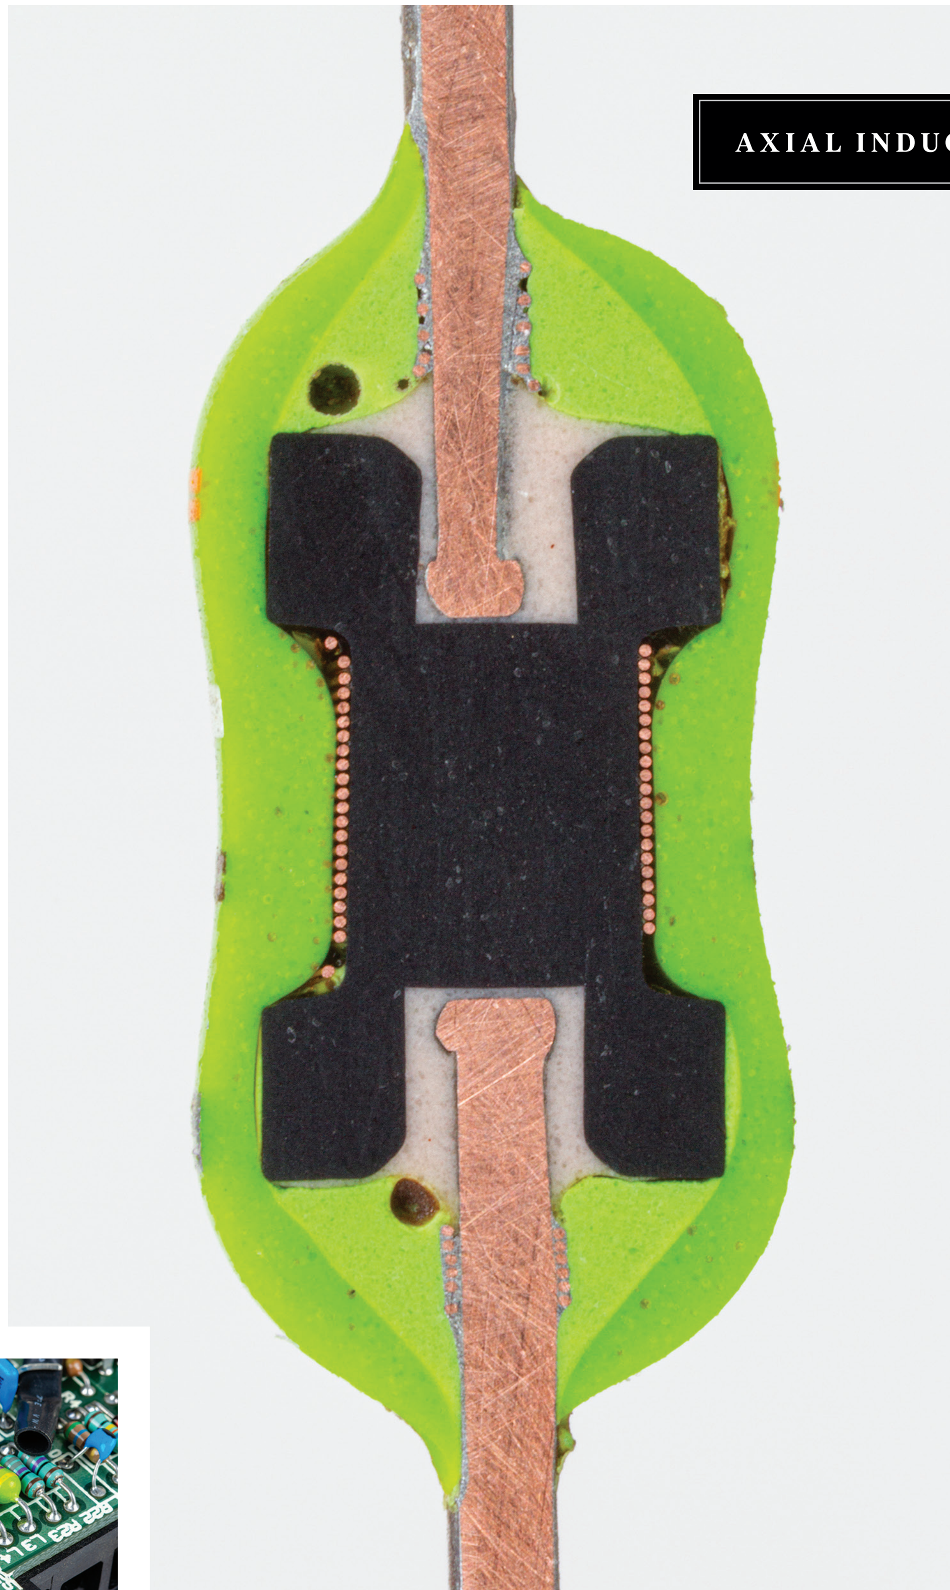

The Panopticon v. the Capitol Rioters

CONTINUED FROM P. 37

Lynch points out that geofence warrants were also used to investigate possible arsons that occurred during protests over police brutality in Seattle, in 2020. Even though the fires were set at a known location at a known time, the warrants sought location data for all devices on an entire city block over a 75-minute period, during a Black Lives Matter protest. "I think that we would all agree that [the protest] was constitutionally protected First Amendment activity," she says. "That information should never be in the hands of law enforcement, because it chills people from feeling comfortable speaking out against the government."

Google told *Spectrum* that it examines all geofence warrants closely for legal validity and constitutional concerns. It says it routinely pushes back on overbroad demands, and in some cases refuses to produce any information at all.

Geofences target places, not people—and that's a problem

OF COURSE, the idea of staking out a particular area for scrutiny is old hat. "Look at every car parked on Elm Street," says the detective, in just about any procedural, ever. What's new is the ability to survey any area immediately, easily, and over a wide range of databases—every phone call placed, car parked, person employed, credit-card transaction made, and pizza sold.

And indeed, the high-tech investigations around the Capitol breach went far beyond suspects' phones to include Uber rides, users' search history, Apple iCloud, and Amazon. The FBI noted that one suspect, Hatchet Speed, a U.S. Navy reserve officer assigned to the U.S. National Reconnaissance Office, had purchased a black face mask and black "Samurai Tactical Wakizashi Tactical" backpack on Amazon, both of which he was seen wearing in Capitol CCTV footage on 6 January. Speed was arrested in June 2022.

Unsurprisingly, after the deadly riot, some of those present deleted their social-media posts, pictures, and accounts. One suspect threw his phone into the Atlantic Ocean. Annie Howell of Swoyersville, Pa., allegedly posted videos of her clashes inside the Capitol with law enforcement. According to her charging document, on 26 January 2021, Howell conducted a factory reset of her Apple iPhone, without backing up data from her online iCloud account. In a Facebook conversation with her father from her computer, he told her, "Stay off the clouds! They are how they are screwing with us."

The legal documents allege that around 150 others also attempted to delete data and accounts. For many, it was far too late. "The FBI's really good at finding information that's deleted, because, as you might know, if you delete a text or an app on a cellphone, it's not really deleted," an FBI agent told a 6 January suspect during an interrogation, as reported in one court filing. Investigators were indeed able to recover chats, social-media posts, call records, photos, videos, and location data from many devices and accounts that suspects thought they had permanently consigned to the digital trash can. The FBI even used such efforts to identify suspects: It asked Google to single out those devices in the geofence warrant whose users had attempted to delete their location history in the days following the siege. That process netted an additional 37 people. In March 2022, Howell was sentenced to 60 days in jail.

Raising a hue and cry—digitally

PERHAPS THE BIGGEST innovation in the 6 January investigations was nothing that law enforcement itself did, but rather the general public's response. Using tools and processes pioneered by open-source investigation organizations like Bellingcat, websites such as Jan6attack.com and Sedition Hunters provided a forum for ordinary people in the United States and around the world to analyze and speculate (sometimes correctly, sometimes wrongly) on the identity of rioters. The FBI cited such efforts in 63 legal documents.

Nonprofit investigative newsroom ProPublica became involved when a source provided 30 terabytes of video—over a million video clips—that had been scraped from the social-media network Parler. "One thing that was really helpful was that Parler wasn't built very well," says Al Shaw, deputy editor on ProPublica's News Application Team. "There was all this metadata still attached to the files when they were leaked. We had geo information, what cellphone they were using, time stamps, and a bunch of other data."

ProPublica filtered the videos by geolocation and other metadata, but soon realized that not all the data was accurate. So journalists went through videos manually to check that those that appeared to have been shot inside the Capitol actually were. ProPublica ended up with 2,500 videos that it could definitively place in the Senate complex on 6 January.

It quickly published 500 of these videos online. Scrolling through the videos is like fast-forwarding through that chaotic day all over again. "One of the design ideas was, can we build a 'sad TikTok'?" says Shaw. "It's got a similar interface to TikTok or Instagram, where you're seeing what's going on generally in chronological order." ProPublica's videos were cited by the DOJ in at least 24 cases.

The remaining 2,000 Parler videos shot from 6 January are now languishing on ProPublica's servers and could almost certainly help identify more rioters. And the hundreds of thousands of videos discarded in the filtering process could very well contain evidence of further crimes and misdemeanors, as could the thousands of unsearched smartphones and unscraped social-media accounts of other people who went to Washington that day.

But at some point, says EFF's Lynch, we should ask what we're really fighting for. "We could, of course, solve more crime if we let police into everybody's house," she says. "But that's not the way our country is set up, and if we want to maintain a democracy, there have to be limits on surveillance technologies. The technology has advanced faster than the law can keep up."

In practice, that means that some federal courts have found geofence warrants unconstitutional, while others continue to permit their use. Similarly, some jurisdictions are limiting the retention of ALPR data by law-enforcement agencies and the use of facial-recognition technologies by police. Meanwhile, though, private companies are mining ever more open-source images and location information for profit.

In the eternal struggle between security and privacy, the best that digital-rights activists can hope for is to watch the investigators as closely as they are watching us. ■

The variations of *Pong* created by the OpenAI Codex vary widely in ball and paddle size and color and how scores are displayed. Sometimes the code results in an unplayable game, such as at the bottom right, where the player paddles have been placed on top of each other.

An Infinity of *Pong*

CONTINUED FROM P. 17

structure common to many video games, namely a list of libraries we'd like to use, a call to process events (such as keypresses), a call to update the game state based on those events, and a call to display the updated state on the screen.

How to use those libraries and fill out the calls is up to the AI. The key to turning this generic structure into a *Pong* game are the embedded comments—optional in source code written by humans, really useful in prompts. The comments describe the gameplay in plain English—for example, "The game includes the following classes...Ball: This class represents the ball. It has a position, a velocity, and a debug attributes [*sic*]. *Pong*: This class represents the game itself. It has two paddles and a ball. It knows how to check when the game is over." (You can play *Pong* games with the Raspberry Pi Pico W at Hackaday.io; my container and prompt code are on the site.)

What comes back from the AI is about 300 lines of code. In my early attempts the code would fail to display the game because the version of the MicroPython framebuffer library that works with my module is different from the framebuffer libraries the OpenAI Codex was trained on. The solution was to add the descriptions of the methods my library uses as prompt comments, for example: "def rectangle(self, x, y, w, h, c)." Another issue was that many of the training examples used global variables, whereas my initial prompt defined variables as attributes scoped to live inside individual classes, which is generally a better practice. I eventually had to give up, go with the flow, and declare my variables as global.

The code that comes back from my current prompt produces a workable *Pong* game about 80 percent of the time. Sometimes the game doesn't work at all, and sometimes it produces something that runs but isn't quite *Pong*, such as when it allows the paddles to be moved left and right in addition to up and down. Sometimes it's two human players, and other times you play against the machine. Since it is not specified in the prompt, Codex takes either of the two options. When you play against the machine, it's always interesting to see how Codex has implemented that part of code logic.

So who is the author of this code? Certainly there are legal disputes stemming from, for example, how this code should be licensed, as much of the training set is based on open-source software that imposes specific licensing conditions on code derived from it. But licenses and ownership are separate from authorship, and with regard to the latter I believe it belongs to the programmer who uses the AI tool and verifies the results, as would be the case if you created artwork with a painting program made by a company and used their brushes and filters.

As for my project, the next step is to look at more complex games. The 1986 arcade hit *Arkanoid* on demand, anyone? ■

Illustration by James Provost

BY ALLISON MARSH

Past Forward

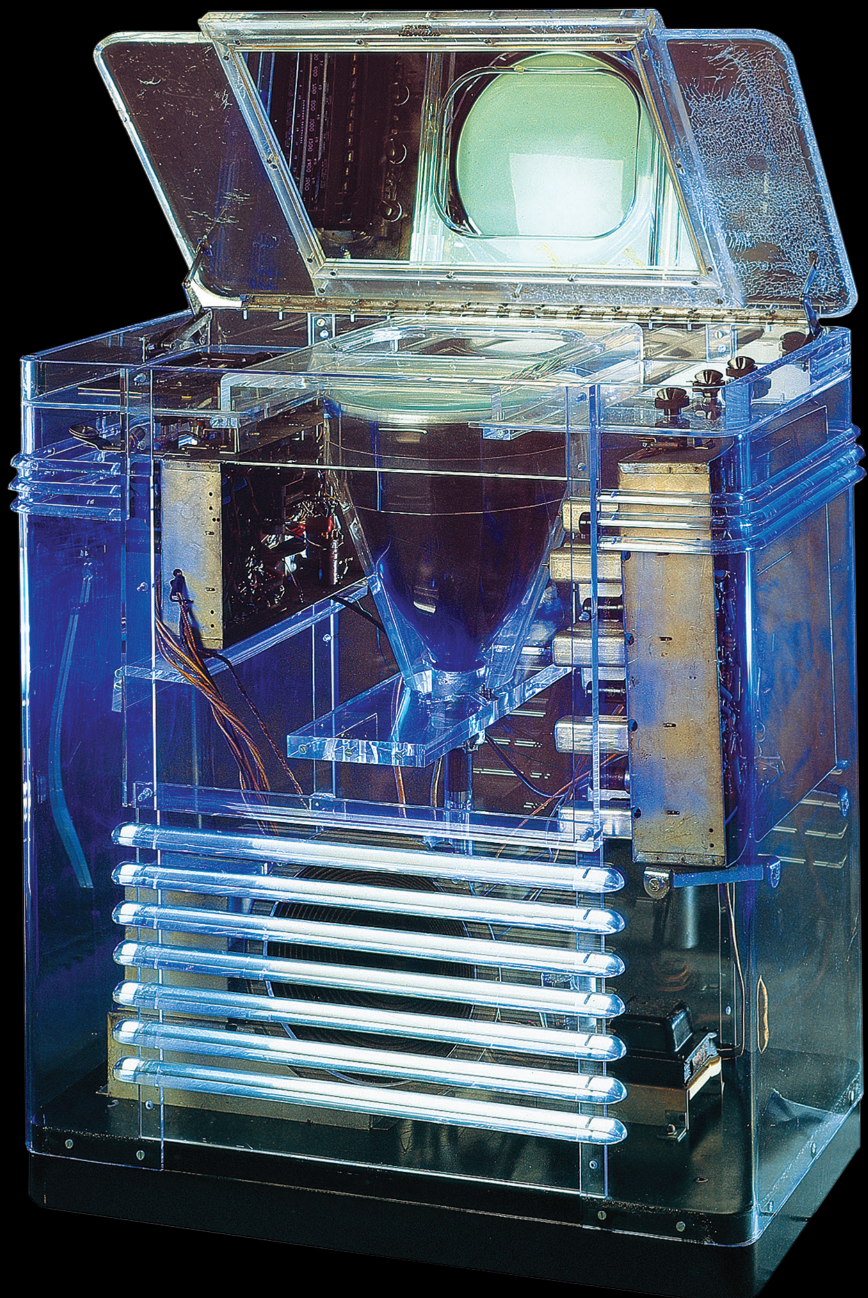

Yesterday's TV of Tomorrow

The 1939 New York World's Fair promised visitors a glimpse into the "World of Tomorrow." At the RCA pavilion, that future included a Lucite-encased television. The clear cabinetry of the TRK-12 Phantom Teleceiver allowed curious spectators to see how the vertically positioned cathode-ray tube projected a live broadcast image onto a 30.5-centimeter (12-inch) mirror in the cabinet lid. But the TRK-12 wasn't just a television—it was the first multimedia center. Standing 102 cm tall and weighing more than 91 kilograms, the console combined a TV set, three-band radio, and optional Victrola phonograph. It was available for purchase for US $600—nearly $13,000 today—albeit with a tamer walnut veneer cabinet. ■

FOR MORE ON THE HISTORY OF THE TRK-12, see spectrum.ieee.org/pastforward-feb2023

MIA KLEIN/MZTV MUSEUM OF TELEVISION