Fixing AI's Zero Problem ›
Sparsity lets AI models grow without wasting energy
P.28

Inside the Living Heart ›
Virtual twins cut the risk of complex procedures
P.34

Return to Trinity ›
Forgotten photos capture the first atom-bomb test
P.40

FOR THE
TECHNOLOGY
INSIDER
JUNE 2026

IEEE Spectrum

SARDINIA

The island's energy future collides with its turbulent past

VOLUME 63 / ISSUE 6 JUNE 2026

Contents

LOS ALAMOS NATIONAL LABORATORY

ON THE COVER:
Photo by Luigi Avantaggiato

BY HARRY GOLDSTEIN

The Roots of Renewables Resistance

To see why Sardinians reject renewables, you need to understand 2,700 years of invasion and exploitation

The phrase "not in my backyard" is the rallying cry of citizens everywhere resisting projects proposed for their locality. Whether it's affordable housing, a waste treatment plant, or a new data center, they may recognize the benefit of the activity. They just don't want it near them. And the roots of that resistance differ from place to place. When it comes to the ongoing transition from fossil fuels to renewables, companies and policymakers need to know where, exactly, people are coming from.

The Italian island of Sardinia is a textbook example. As *IEEE Spectrum*'s power and energy editor Emily Waltz discovered when she traveled there last October, Sardinian opposition to wind and solar projects runs deep. It spurred a quarter of the voting population to queue up in public squares in 2024 to sign a petition banning all construction of renewable energy.

Waltz was surprised. She went there to see a promising new grid-scale energy storage system that uses domes inflated with carbon dioxide. While reporting on that project, she interviewed residents, engineers, activists, and professors about their attitudes toward climate change and the Italian government's grand plans for renewable energy on the island. And Waltz soon learned of Sardinians' profound antipathy toward renewable energy and its deep ties to a history of invasion, occupation, and exploitation stretching back 2,700 years.

It started with the Phoenicians and then extended through the Romans, the Byzantines, and the Iberians. Sardinia was absorbed into a newly unified Italy in 1861, and it became an autonomous region of Italy in 1948. The island's population is justifiably suspicious of outsiders, including the Italian government.

Energy editor Emily Waltz [left] toured an old Sardinian coal mine, nearly 500 meters down, that's being converted into a data center powered by a pumped hydro and battery system. Her guide was Energy Vault executive Luca Manzella [next to her].

"When you're in Sardinia, the weight of history—you can feel it in the air," Waltz told me. "And it gets passed down from one generation to the next."

Now, Italy needs Sardinia to produce even more power to meet the country's climate goals—something that Sardinians see as Rome's problem, not theirs. "Sardinia already exports about 30 percent of its electricity. It's not like they need more," Waltz says. "So it's hard to make the case to build, build, build."

The result of Waltz's old-fashioned shoe-leather reporting is this month's cover story [p. 20]. She notes that the Sardinians she talked to aren't climate-change deniers, and they don't object to renewables per se. They just don't like the way corporations and Italian policymakers are trying to plug into Sardinia like it's one giant battery rather than home to an ancient and proud people.

"I think Sardinians would be more receptive to renewable projects if it was more of a ground-up, grassroots approach," Waltz says. Indeed, this homegrown effort is already working in some places in Sardinia. She knows of more than 50 projects, called energy communities, where the residents are deploying renewables themselves. The idea also holds promise for other places struggling to get locals to buy into the renewable-energy transition.

The Sardinian experience is both a cautionary tale and a blueprint. Ignore the weight of history that communities carry and your project risks failure. Meet the people where they are and you might just get somewhere. The same lesson applies whether you're in Sulawesi or sub-Saharan Africa. You just have to show up to learn it. ■

CORRECTION: In "Rise of the Autonomous Attack Drones" [April 2026], we referred to certain militants in Africa as "Islamic terrorists." We regret the use of that phrase, which wrongly characterizes an entire religion.

PORTRAIT BY SERGIO ALBIAC; EMILY WALTZ

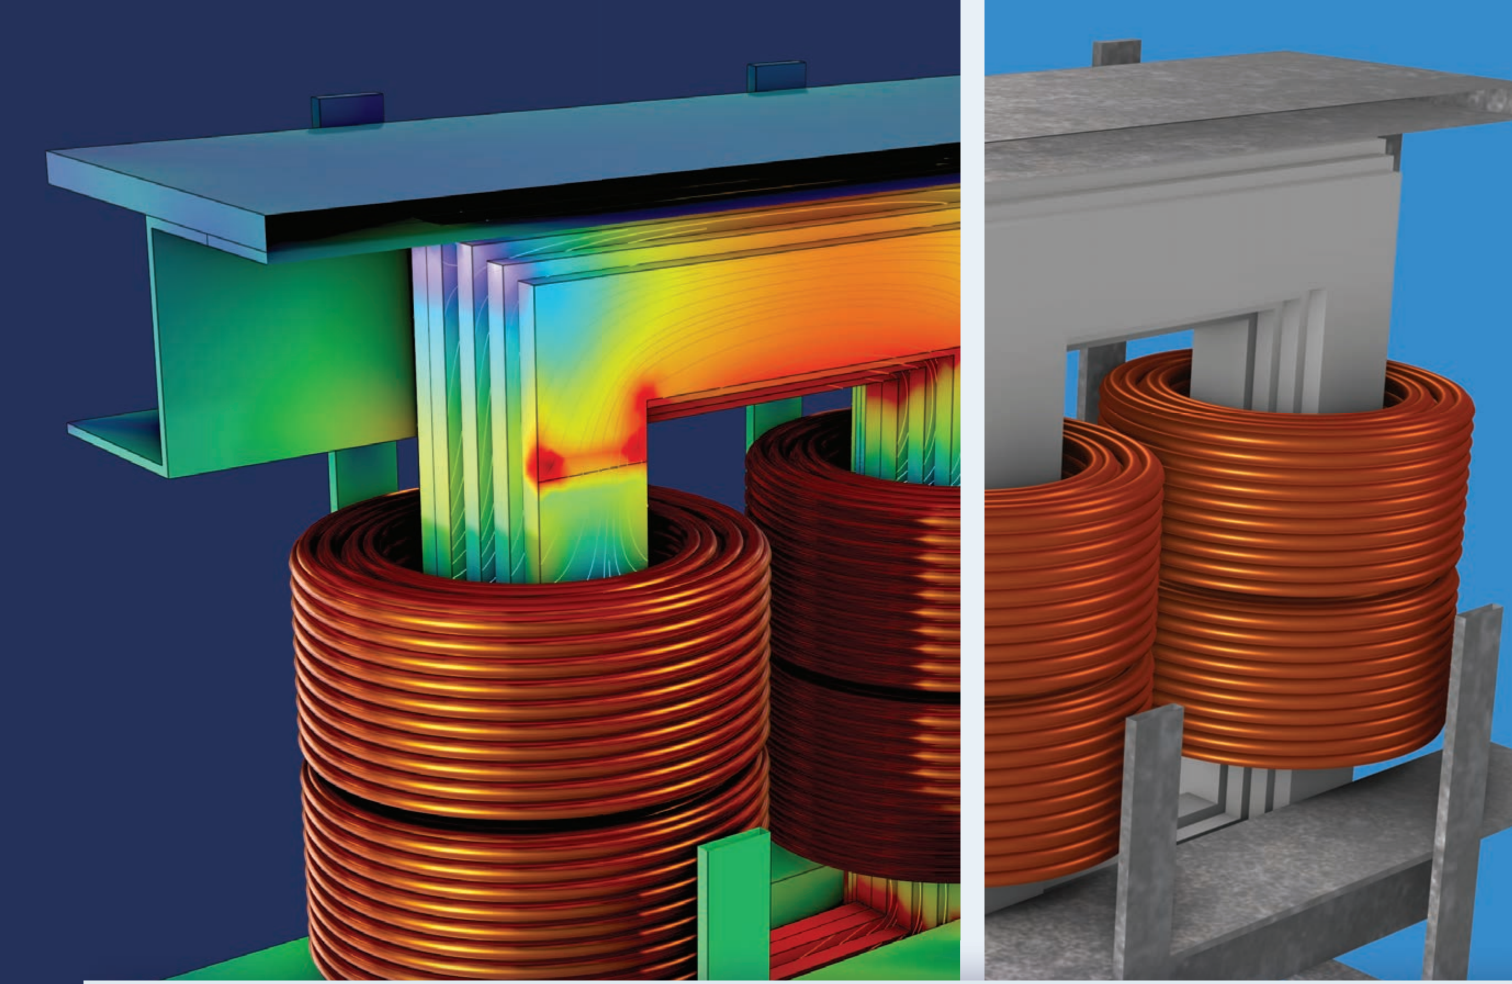

CONTRIBUTORS

OLIVIA HSU & KALHAN KOUL

Hsu is a postdoctoral researcher at Stanford who will soon join the faculty at Carnegie Mellon as an assistant professor of electrical and computer engineering. In this issue, Hsu and her coauthor Koul, formerly a graduate student at Stanford and now a senior GPU architect at Nvidia, describe how they're working to quell AI's insatiable energy appetite [p. 28]. By building hardware that's wired to bypass zeros in unnecessary computations, they hope to empower both software and hardware architects to come up with more creative, waste-free solutions.

STEVE LEVINE

Levine is senior director of virtual human modeling at Dassault Systèmes, based in France. In his feature on creating a dynamic 3D model, or virtual twin, of the human body [p. 34], he describes how his daughter's health condition inspired him to start the Living Heart Project and the progress it has made toward personalized, predictive medicine. "Watching my daughter's doctors guess about her unique heart, I realized modern physicians need the same access to physics that engineers have," he says. "We've brought that thinking back to medicine."

EMILY SEYL

Seyl is the author of *Trinity: An Illustrated History of the World's First Atomic Test*, which we excerpt in this issue ("The Last Human Will See What We Saw," p. 40). She is also a writer at the National Security Research Center at Los Alamos National Laboratory. Although she did not train as an engineer, she finds herself drawn to stories of engineering. "My grandfather is an electrical engineer," she says. "He worked at NASA during the space race. So I guess you could say it's in my blood."

JACKIE SNOW

Snow is a Los Angeles–based freelance journalist covering AI, with bylines in *The New York Times, The Wall Street Journal, National Geographic*, and Quartz. In this issue, she reports on the machine learning tools being used to manage the dwindling Colorado River [p. 5]. "I usually report on AI draining water, not helping to manage it," she says. "The inversion was a welcome change, especially because some of that Colorado River water comes out of my own tap."

IEEE Spectrum

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
DIRECTOR OF DIGITAL INNOVATION Erico Guizzo, e.guizzo@ieee.org
EDITORIAL DIRECTOR, CONTENT DEVELOPMENT
Glenn Zorpette, g.zorpette@ieee.org

FEATURES EDITOR Eliza Strickland, e.strickland@ieee.org

SPECIAL PROJECTS DIRECTOR Stephen Cass, cass.s@ieee.org

SENIOR EDITORS
Evan Ackerman, evan.ackerman@ieee.org
Lucas Laursen, l.laursen@ieee.org
Samuel K. Moore, s.k.moore@ieee.org

NEWS MANAGER Michael Koziol, m.koziol@ieee.org

SENIOR ASSOCIATE EDITORS
Margo Anderson, m.k.anderson@ieee.org
Emily Waltz, e.waltz@ieee.org

ASSOCIATE EDITORS
Dina Genkina, d.genkina@ieee.org
Willie D. Jones, w.jones@ieee.org

ASSISTANT EDITOR Gwendolyn Rak, g.rak@ieee.org

PRODUCT AND AUDIENCE DEVELOPMENT DIRECTOR
Tim Warder, t.warder@ieee.org
PRODUCT MANAGER Priyanka Shitole, p.shitole@ieee.org
AUDIENCE DEVELOPMENT MANAGER
Laura Bridgeman, l.bridgeman@ieee.org
COMMUNITY MANAGER
Kohava Mendelsohn, k.mendelsohn@ieee.org
DIGITAL PRODUCT DESIGNER Erik Vrielink, e.vrielink@ieee.org

ART & PRODUCTION
ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSISTANT PHOTO EDITOR Nicole Millman, n.millman@ieee.org
PRINT PRODUCTION SPECIALIST
Sylvana Meneses, s.meneses@ieee.org
DIGITAL PRODUCTION MANAGER
Michael Spector, m.spector@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
RESEARCH COPY EDITOR Cheryl Della Pietra, c.dellapietra@ieee.org

ADMINISTRATOR Jose Zambrano, jose.zambrano@ieee.org

CONTRIBUTING EDITORS Rina Diane Caballar, Robert N. Charette, Charles Q. Choi, Tom Clynes, Peter Fairley, Edd Gent, W. Wayt Gibbs, Mark Harris, Allison Marsh, Julianne Pepitone, Ned Potter, Tereza Pultarova, Matthew S. Smith, Lawrence Ulrich

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

SUBSCRIPTION SERVICES
IEEE Members member-services@ieee.org
Non-members +1 866 363 2304

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org

ADVERTISING PRODUCTION +1 732 562 6334

EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM*
Harry Goldstein, *Chair*; Muhammad Alaa, Karanveer Anand, Ella M. Atkins, Sangyeun Cho, Hugh Durrant-Whyte, Benjamin Gross, Carlos Gutierrez, Lawrence O. Hall, Benjamin Kroposki, Somdeb Majumdar, Khaled Mokhtar, Ramune Nagisetty, Paul Nielsen, Ketra Schmidt, Gaurav Sharma, Fan Shi, Mini S. Thomas, Andrej Zdravkovic, Lidong Zhou

EDITORIAL ADVISORY BOARD, *THE INSTITUTE*
Kathy Pretz, *Chair*; Stamatis Dragoumanos, Madeleine Glick, Harry Goldstein, Lawrence O. Hall, Francesca Iacopi, Sai Prashanth Mallellu, Shashi Raj Pandey, John Purvis, Chenyang Xu

EDITORIAL CORRESPONDENCE
IEEE Spectrum, 3 Park Ave., 17th Floor,
New York, NY 10016-5997 Tel: +1 212 419 7555

ADVERTISING INQUIRIES WILEY
Vanessa Winde +49 620 160 6721, vwinde@wiley.com
Chris Darch +1 206 867 4165, cdarch@wiley.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.*

COPYRIGHTS AND TRADEMARKS *IEEE Spectrum* is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Responsibility for the substance of articles rests upon the authors, not IEEE, its organizational units, or its members. Articles do not represent official positions of IEEE. Readers may post comments online; comments may be excerpted for publication. IEEE reserves the right to reject any advertising.

IEEE

IEEE BOARD OF DIRECTORS
PRESIDENT&CEO Mary Ellen Randall, president@ieee.org
+1 732 562 3928 Fax: +1 732 981 9515
PRESIDENT-ELECT Jill I. Gostin
PAST PRESIDENT Kathleen A. Kramer
TREASURER Gerardo Barbosa
SECRETARY Jenifer Castillo Rodriguez

VICE PRESIDENTS
Timothy P. Kurzweg, *Educational Activities;* Andrew D. Lowery, *Member & Geographic Activities;* W. Clem Karl, *Publication Services & Products;* F.D. Tan, *Technical Activities;* Gary R. Hoffman, *President, Standards Association;* Barry C. Tilton, *President, IEEE-USA*

DIVISION DIRECTORS
Fernando J. Guarin (I); H. Alan Mantooth (II); Sarah Kate Wilson (III); Charles M. Jackson (IV); Leila De Floriani (V); Russell D. Meier (VI); Noel M. Schulz (VII); Cecilia Metra (VIII); Cristina Olaverri-Monreal (IX); Ravinder S. Dahiya (X)

REGION DIRECTORS
Charles P. Rubenstein (1); Felicia Reinhart (2); Sonya L. Dillard (3); Constance A. Kelly (4); Christopher B. Sanderson (5); Joseph C. Wei (6); Wahab Almuhtadi (7); Dusanka Boskovic (8); Jose-Ignacio Castillo-Velazquez (9); Takako Hashimoto (10)

IEEE STAFF
EXECUTIVE DIRECTOR & COO Sophia A. Muirhead
+1 732 562 5400, s.muirhead@ieee.org
CHIEF FINANCIAL OFFICER Kelly Armstrong
+1 732 562 6011, k.armstrong@ieee.org
MANAGING DIRECTOR,TECHNICAL ACTIVITIES Mojdeh Bahar
+1 732 562 3850, m.bahar@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Ahsaki Benion +1 212 705 8975, a.benion@ieee.org
MANAGING DIRECTOR, PUBLICATIONS Steven Heffner
+1 212 705 8958, s.heffner@ieee.org
CHIEF GOVERNANCE OFFICER Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MANAGING DIRECTOR, CONFERENCES, EVENTS AND EXPERIENCES Marie Hunter
+ 1 732 465 5889, m.m.hunter@ieee.org
MANAGING DIRECTOR, MEMBER & GEOGRAPHIC ACTIVITIES
Cecelia Jankowski +1 732 562 5504, c.jankowski@ieee.org
MANAGING DIRECTOR, EDUCATIONAL ACTIVITIES
Jamie Moesch +1 732 562 5514, j.moesch@ieee.org
CHIEF MARKETING OFFICER Jayne O'Brien
+1 732 981 3433, jayne.obrien@ieee.org
MANAGING DIRECTOR, IEEE STANDARDS Alpesh Shah
+1 732 465 6467, alpesh.shah@ieee.org
CHIEF INFORMATION AND DIGITAL OFFICER Jeff Strohschein
+1 732 562 5332, j.strohschein@ieee.org
CHIEF OF STAFF Allison Taylor
+1 732 562 6068, allison.taylor@ieee.org
CHIEF HUMAN RESOURCES OFFICER Cheri N. Collins Wideman
+1 732 562 5335, c.wideman@ieee.org
MANAGING DIRECTOR, IEEE-USA Russell T. Harrison
+1 202 530 8326, r.t.harrison@ieee.org

IEEE PUBLICATION SERVICES & PRODUCTS BOARD
W. Clem Karl, *Chair*; John T. Barr, Nazanin Bassiri-Gharb, Sergio Benedetto, Eddie Custovic, Stephen Dukes, Alessandro Foi, Lawrence O. Hall, Charles Hansen, Warunika Hippola, James Irvine, Michael Jensen, Athanasios Kakarountas, James Keller, Imran Mehdi, Khaled Mokhtar, Paolo Montuschi, Zorana Popovic, Mehrdad Saif, Gianluca Setti, Gaurav Sharma, Sarah Spurgeon, Jaideep Vaidya, Maria Elena Valcher, Patrick Wheeler, Steve Yurkovich

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. © 2026 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 63, Issue No. 6. 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: United States and Canada: $74.95. Elsewhere: $99.95. 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 3401 N. Heartland Dr., Liberty, MO 64068, U.S.A. Printed at Spenta Multimedia Pvt. Ltd., Plot 15, 16 & 21/1, Village Chikhloli, Morivali, MIDC, Ambernath (West), Dist. Thane. *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

The Colorado River Basin extends across seven U.S. states, including Arizona, where the river meanders through Horseshoe Bend. The river system's dwindling water supply is forcing states to make difficult decisions.

SOURCE IMAGES: WATTS; USDA NATURAL RESOURCES CONSERVATION SERVICE

ARTIFICIAL INTELLIGENCE

AI Reveals the Fraught Future of the Colorado River

› Simulations show the trade-offs as states fight over water

BY JACKIE SNOW

The Colorado River begins as snow. Every spring, the mountain snowpack of the Rockies melts into streams that then feed into reservoirs supplying 40 million people across seven U.S. states. The system has worked, more or less, for a century. That century is over.

By some measures, 2026 is shaping up to be the worst year the river has seen since records began. Flows are down 20 percent from 2000 levels. Lake Powell, the reservoir straddling Utah and Arizona, may drop below the

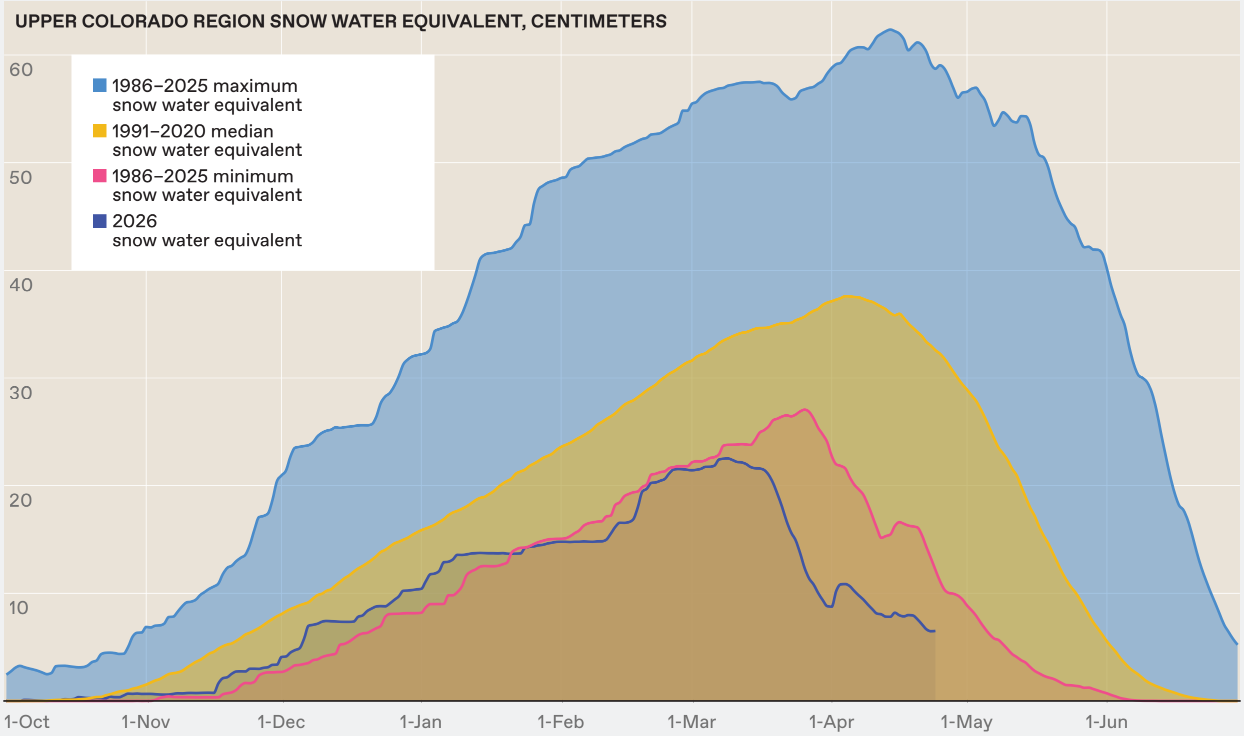

Snow water equivalent measures how much liquid water is contained within a snowpack when it melts. A dry winter [December to February in the Upper Colorado River Basin] combined with a warm spring has meant a dangerously low snowpack feeding into the Colorado River this year.

SOURCE: USDA NATURAL RESOURCES CONSERVATION SERVICE

threshold for generating hydropower before the year is out. The negotiations among the seven states over how to share what's left have collapsed twice, and the U.S. federal government has implemented emergency plans to release water from Flaming Gorge Reservoir on the Utah-Wyoming border to keep Lake Powell's water level up.

While the states argue and the river shrinks, a growing set of machine learning tools is being deployed across the basin. Federal water managers are running millions of simulations to stress-test reservoir strategies against different possible futures. Researchers are forecasting streamflow months out by using satellite data and deep learning. These technologies don't promise to resolve the crisis, but they're making the trade-offs visible. They're showing, more precisely than ever before, what each decision will cost.

Nobody manages more of the Colorado River's daily operations than the U.S. Bureau of Reclamation, which is responsible for imposing any water-sharing plan from the federal government. The bureau also makes decisions about how much water flows from Lake Powell and Lake Mead, the two largest reservoirs in the country.

The agency is not new to sophisticated computer modeling. Machine learning tools are adding to Reclamation's tool kit, says Chris Frans, the bureau's water-availability research coordinator, and they are already informing real operational decisions.

The clearest gains are in forecasting the flow of streams that feed into the river. Machine learning techniques—using data from satellites and weather stations well outside the basin—now outperform traditional methods across a range of conditions. Forecasts update every hour. In some areas, managers are getting five to seven days of advance warning on flood events, compared with three in the past, which gives them time to reduce the water in reservoirs before high inflows arrive.

The scale of scenario modeling, in which possible futures are constructed and analyzed, has also expanded dramatically. A decade ago, running 100,000 simulations of possible futures was a landmark study. Now, says Alan Butler, who manages Reclamation's research and modeling group for the Lower Colorado River Basin, millions of simulations feed the analytical tools being used to create the new guidelines for water usage. Those simulations map out how different operating strategies perform across widely varying futures—making the compromises between them harder to ignore.

Knowing how much water is coming is one problem. Deciding who gets it is another. At the center of that process is the Colorado River Simulation System (CRSS), which models how water moves through the basin's reservoirs, canals, and pipelines under more than a century of legal and regulatory constraints. This Reclamation model is an imperfect representation, but it has been the foundation of river negotiations for decades.

A tool called RiverWare, first developed in the early 1990s with the University of Colorado Boulder and the Tennessee Valley Authority, lets states, cities, and Native American governments run their own scenarios through CRSS. Before that, these groups didn't have confidence in Reclamation's numbers. "There was just this huge lack of trust," says Edith Zagona, a Boulder professor who directs the Center for Advanced Decision Support for Water and Environmental Systems, which built RiverWare. The solution was letting stakeholders inspect the assumptions built into the RiverWare model—how much water was available, how it could be used, and under what rules.

Getting regional authorities to trust the model turned out to be the easier problem. The harder one is what to do when CRSS can't say which future is most probable because climate change makes past data unreliable for prediction. That issue drove Zagona toward what's called decision-making under deep uncertainty, which trades prediction for stress-testing policies against thousands of possible futures.

Zagona's group developed a web-based tool called Borg-RiverWare with Reclamation and the consulting firm Virga Labs to put the new strategy into practice. The tool runs CRSS across more than 8,000 possible future water-supply scenarios to show how different management strategies hold up against the full range of what climate change might bring. At its center is an evolutionary algorithm called Borg, which generates and iteratively refines those strategies, searching for plans that perform well across many scenarios. The result is a set of trade-offs, not a single answer.

Borg-RiverWare has already shaped the ongoing negotiations over the river's next set of operating rules, giving Reclamation and the jurisdictions a common analytical foundation for evaluating proposals. Those tools give stakeholders a common analytical foundation for negotiations. Now Zagona's center is pushing the approach further. A system in development would let negotiating parties test competing proposals on the fly, showing how one side's policy choices would ripple through the system and identifying areas of potential compromise during the negotiation itself.

Reclamation and Zagona's center aren't the only ones trying to see further into the river's future. At Metropolitan State University of Denver, a team led by Mohammad Valipour has been building a forecasting system that uses deep learning to issue drought warnings across seven rivers in Colorado, from seven days to six months out. The goal, Valipour says, is a statewide drought-alarm system that gives farmers and water managers more time to respond.

At Utah State University, Soukaina Filali Boubrahimi is attacking a different problem: how conditions at one point in the river ripple downstream weeks later. Using a graph neural network that treats each monitoring station as a node, her team built a map of the river's interdependencies. If it works for the Colorado River, one of the most contested water systems in the world, the approach could be extended to other overtaxed basins, she says.

"If you can figure out the Colorado River," says Filali Boubrahimi, "anyone else dealing with a stressed river system is going to be interested in what you learned."

Peering into the river's future is hampered by using historical data that describe a river that no longer exists. Valipour found that feeding his models only the last decade of data outperformed using longer records. Filali Boubrahimi's model struggles most in drought conditions, because recent prolonged droughts don't resemble the historical training data. One workaround is to train models on data from basins that have already experienced what may be in the Colorado River's future.

Even so, better forecasts do not resolve the central problem: Who should bear the cost of a drier future. The cuts coming to the basin are going to be enormous and they will fall mostly on agriculture, says Brad Udall, a water and climate research scientist at Colorado State University's Colorado Water Center. They may fundamentally reshape communities that have built their economies around water for generations.

The tools, by most measures, are doing exactly what they were built to do: The negotiating parties understand what is coming, and they are not disputing the projections. Zagona, who has worked on the Colorado River for 45 years, sees reasons for optimism. "The tools are bringing people to the table," she says. "They're at the table arguing. But at least they're at the table." ■

TRANSPORTATION

Optical Fibers Sense Risks on the Rails

MONITORING THE SAFETY of today's vast rail networks is a huge challenge. Researchers in China might have a way to make it a lot easier. They propose using existing fiber-optic communication cables buried underground alongside railway tracks to spot potential trouble.

Their approach uses distributed acoustic sensing, a technique in which pulses of laser light are sent along the cable. When something nearby—such as a passing train—causes the fiber to vibrate, a portion of the light is scattered back along the cable.

In their study, the researchers from Nanjing Fiber Photonics Technology, Nanjing University, and Southeast University, also in Nanjing, trained machine learning models to identify the sources of different scattering patterns.

The models succeeded at determining a range of railway safety indicators, including train trajectories, defects in train wheels or their connecting axles, and broken acoustic barriers alongside the railway tracks. The system could also detect humans climbing over trackside fences, rocks falling on the track, and nearby construction activity.

The results were published in March in the *Journal of Optical Communications and Networking*.

—Michelle Hampson

BIOMEDICAL

“Nanodrums” Identify Bouncing Bacteria

Each germ’s movements generate a sonic fingerprint to aid diagnosis

BY RAHUL RAO

Researchers at Delft University of Technology, in the Netherlands, have crafted a nanoscale “drum kit” that can detect a bacterium’s movements. They’ve discovered that different species of bacteria move in different ways, making it possible to identify a bacterium, and diagnose an infection, from the “song” a bacterium’s motion causes.

“We can really look at the level of a single cell,” says Farbod Alijani, a mechanical engineer at TU Delft. “We have that sensitivity.”

> “It’s very noisy, like a wind tunnel.”
>
> —ALEKSANDRE JAPARIDZE, SOUNDCELL

A researcher deposits antibiotics onto a bacteria sample to test for resistance using SoundCell’s nanodrum platform.

SOUNDCELL

The TU Delft researchers call their miniature instrument a “nanodrum.” Its drumhead is fashioned from two graphene sheets, each less than 1 nanometer thick, laid atop an 8-micrometer-wide cavity. This size fits most bacteria, which are about 1 to 10 micrometers long. Alijani and his colleagues reported their latest findings this March in *ACS Sensors*.

Starting in 2019, the researchers noticed that, if a living bacterium settled on the graphene sheet, it would beat a pattern on the drumhead. They were detecting the life-form’s subtle motions, such as the whirling of the propeller-like flagellum a bacterium uses to move about. When the drumhead moved, it left signals on a beam of laser light reflected off the surface. The drum’s tiny size is key to pinpointing individual bacteria. The Delft researchers were not the first to capture bacteria in motion, but older methods usually had to average the movements of an ensemble of many bacteria tapping on the same surface. By comparison, each graphene drumhead is small enough to isolate—and record—a single bacterium. Graphene has the perfect mix of flexibility and strength, so the ultrathin surface can bend with each subtle bounce on the drum without breaking.

By converting its drumbeat to a soundtrack, the researchers could literally hear the motions of a living bacterium. “It’s very noisy, like a wind tunnel,” says Aleksandre Japaridze, the chief technical officer of SoundCell, a company spun off of TU Delft to commercialize the research. However, “if you kill it with a drug, it’s immediately very silent, and you don’t hear anything.”

In previously published work, when the researchers pumped an antibiotic onto drums housing *E. coli*, the drums mostly fell quiet within hours. When they did the same to *E. coli* they knew to be antibiotic-resistant, the bacteria lived—and the drumming continued. The group worked to refine their technology’s ability to screen bacteria for antibiotic resistance.

But their work took an unexpected turn after an attendee at a conference in 2023 asked Alijani if different bacteria’s motions made different sounds. Unsure of the answer, the researchers wondered how they could find out.

Separating one bacterium from another species by sound alone required a more sophisticated approach than simply determining whether or not a sample was still alive. The team recorded the drumbeats of different infectious bacteria from real patients' samples. Instead of using raw sounds, the researchers processed them into time-frequency spectrograms, which allowed them to more carefully study the patterns of each bacterium's motion.

They trained two different machine learning models to examine a spectrogram and identify its drummer as one of three species: *E. coli*, *Staphylococcus aureus* (responsible for staph infections), or *Klebsiella pneumoniae* (one of the germs that can cause pneumonia).

Both models scored high marks in testing: One classified bacteria with 87 percent accuracy, and the other achieved 88 percent. These results suggest that each species plays different characteristic beats when it moves on the drum.

"It's a completely different way of interpreting the different species," Japaridze says. "Not chemically or biologically, with markers and genes, but just purely on mechanical behavior."

The TU Delft researchers think their drums are a powerful diagnostic tool. Antimicrobial-resistant germs may already be responsible for more than 1 million deaths globally each year and could play a part in millions more. Part of why these bacteria are such potent threats is that the tests for whether a microbe is resistant are relatively slow. Today's tools, which require a patient's sample to be cultured and exposed to different antibiotics, may take days to report whether a microbe is resistant to a given antibiotic. By comparison, nanodrum technology can do this in just an hour or two.

SoundCell was originally spun off to commercialize that ability to determine whether a bacterium is resistant to an antibiotic. But first, the company must show its nanodrums can work in the hospital. SoundCell has repackaged the bulky laboratory system into a smaller device better suited for clinical use, which it has deployed at two hospitals in the Netherlands to verify performance, Japaridze says.

If the tests go well, hospitals in the future may skip the microscope and cultures, instead identifying an infection by listening to the songs of a patient's sample. ■

Transcelestial launched a demo version of its laser-communications terminal into space on board the 6GStarLab mission in November.

TRANSCELESTIAL

TELECOMMUNICATIONS

Can Lasers Replace Radio for Space-to-Ground Comms?

› Getting optical signals through the atmosphere remains a challenge

BY TEREZA PULTAROVA

With the growing number of spacecraft in orbit and the increasing quantities of data beaming back to Earth, radio spectrum is hitting its physical limits. That's why companies and research institutions have been working on higher-bandwidth optical technologies that would remove today's data-transmission bottlenecks.

But using lasers to beam data to Earth comes with challenges that so far have been difficult to solve, such as penetrating cloud cover and pinpointing narrow connections to receivers on the ground. Several companies, however, believe they've finally cracked the problem through a combination of new technology and lessons learned from ground-to-ground optical communications networks.

Perhaps the most ambitious among those companies is Singapore-based Transcelestial, which has been conducting space trials of its space-to-Earth laser-communication terminal after launching a demonstration payload last November. The company has further launches scheduled in 2026 and says it aims to provide fiber-equivalent connectivity from orbit by the end of the decade.

Instead of beaming internet directly to individual users on Earth as SpaceX's Starlink does, Transcelestial envisions delivering tens to hundreds of gigabits to local telecom companies, which would

Transcelestial's optical ground stations [left] are small enough to fit on a building's rooftop, such as this one in Singapore. Transcelestial's space-to-ground laser technology is informed by the company's previous work developing terrestrial optical networks.

TRANSCELESTIAL

further distribute connectivity to users via local ground-based infrastructure. In the future, Rohit Jha, Transcelestial's CEO and cofounder, envisions orbital lasers replacing even undersea cables, offering a cheaper, more reliable service that could not be as easily disrupted by sabotage or natural disasters.

Laser light transmits at higher frequencies than radio waves, and can therefore pack in orders of magnitude more data. The Starlink constellation, which uses a radio link to transmit data back to Earth, offers a peak user bandwidth of about 400 megabits per second—which gets diluted as the number of users in an area grows. By comparison, Transcelestial's test satellite can beam data to Earth at rates up to 1 gigabit per second. The company is aiming for as high as 10 Gb/s for future satellites.

"Scaling for an optical system is actually quite easy," says Jha. "Ultimately, we can deliver 100 Gb/s just by putting more terminals on the satellite. It will be like undersea cables from space."

Lasers are already well established for space-to-space communications in satellite networks. Starlink, for example, has used the tech since 2021 to form an orbital mesh network that can route vast amounts of data through space in real time. The constellation still requires conventional radio signals to return data to Earth, however, so overall network bandwidth remains throttled.

But laser communication tech has made big gains recently. In 2023, NASA tested a record-breaking 200-Gb/s laser link between a ground station and NASA's Pathfinder Technology Demonstrator 3 satellite in low Earth orbit. That same year, the Chinese Academy of Sciences conducted a more modest demo with a 10 Gb/s laser connection between low Earth orbit and the ground. But these experiments required expensive, bespoke equipment, says Mohammad Danesh, Transcelestial's cofounder and chief technology officer. More recently, the Artemis II mission beamed high-definition video back to Earth from near the moon using a 260-Mb/s laser link.

Danesh believes the startup can take the technology mainstream by building on its work commercializing terrestrial point-to-point laser communication systems for internet distribution in hard-to-reach areas on Earth.

"The biggest challenge is building a reliable and scalable optical ground-station network," Danesh says, referring to the stations that communicate with satellites in orbit. "The optical ground stations that people are building today cost millions of dollars, and that's not scalable. You need to be able to manufacture these at scale." What's more, by providing alternative downlink and uplink locations all over the world, a multiple-ground-station system will help overcome the difficulty laser light has getting through clouds, Danesh adds.

Joachim Horwath, the chief technology officer of Germany-based laser communications developer Mynaric, cautions that the challenges presented by clouds and other atmospheric interference might be more difficult to surmount than some think. "Atmospheric conditions remain a key technical challenge," he says, which is why strategies including multiple ground stations or mixing laser and radio communications should still be considered. "Because of this, we don't expect laser communications to replace [radio frequency] entirely."

Lasers, in addition to higher bandwidth, are also much narrower and more focused than radio waves, making them more resilient to jamming and interception.

"You have to be literally within the line of sight of the communication beam to be able to disrupt it," says Laurynas Mačiulis, the CEO of Astrolight. "It's practically very difficult."

Astrolight, based in Lithuania, has also developed a space-to-ground laser-communications terminal, which was launched on 30 March aboard two small satellites developed in Greece by the National and Kapodistrian University of Athens and the Aristotle University of Thessaloniki. The company has a more modest goal than Transcelestial, hoping to enable operators of Earth-observing satellites to retrieve their data faster.

The company previously tested secure laser communication links to transmit data between two ships at sea. In the test, the terminals provided reliable, continuous high-bandwidth communications, even in rainy and foggy weather, for two weeks.

SpaceX is also reportedly looking at space-to-ground laser-communication technology to overcome the bandwidth bottleneck that currently plagues Starlink users in more densely populated areas. ■

ARTIFICIAL INTELLIGENCE

The State of AI in 2026

BY MATTHEW S. SMITH

2026 is a massive year for AI. Anthropic and OpenAI are both aiming for IPOs by the end of the year, headlining an unprecedented amount of investment in the industry.

As it's done every year since 2017, the Stanford Institute for Human-Centered Artificial Intelligence published its AI Index—a comprehensive snapshot of the technology and everything around it. The 2026 edition, coming in at over 400 pages, covers trends in investment, computing capacity, public opinion, and much more. We've compiled four charts that speak to the state of AI in 2026.

For our full selection of 12 charts, see https://spectrum.ieee.org/state-of-ai-index-2026

NUMBER OF NOTABLE AI MODELS

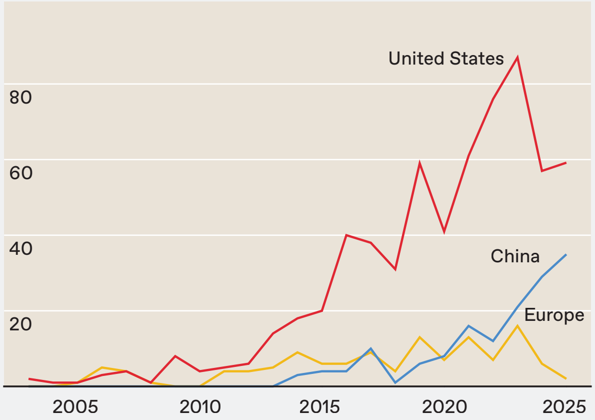

CHINA CLOSES THE GAP: Although the majority of "notable" AI models (meaning ones that have a large number of users or are considered state of the art when they're released) still come from the United States, China has begun to catch up lately. SOURCE: EPOCH AI

CUMULATIVE COMPUTE CAPACITY (H100 EQUIVALENTS, MILLIONS)

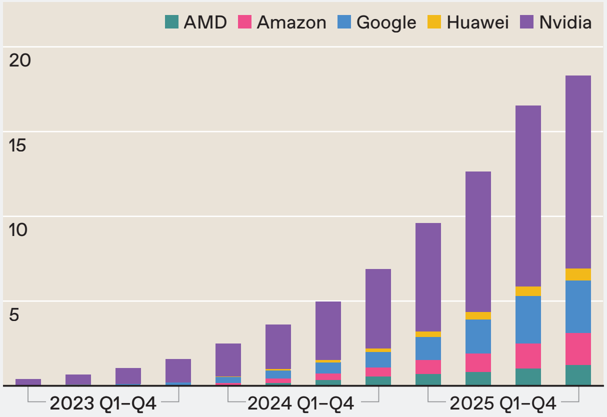

GLOBAL AI COMPUTE SKYROCKETS: Nvidia leads the explosive growth of AI compute capacity being manufactured. According to research from Epoch AI, which quantified compute by using Nvidia's H100 as a measuring stick, total AI compute capacity has increased over 40-fold since early 2023. SOURCE: EPOCH AI

TOTAL INVESTMENT (U.S. DOLLARS, BILLIONS)

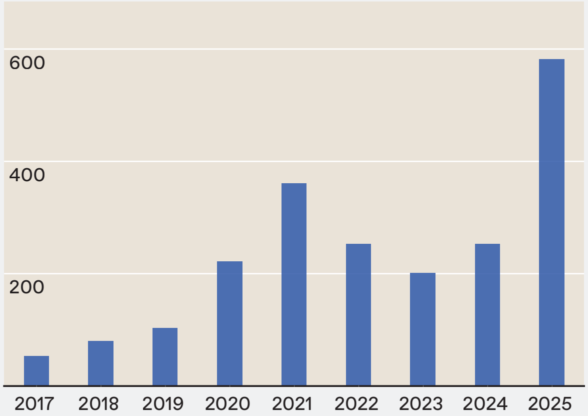

CORPORATE INVESTMENT IS HIGHER THAN EVER: After an earlier peak in 2021, 2025 saw more investment in AI-related companies than ever before—nearly US $600 billion, or more than double the year before. The bulk of that money, roughly $345 billion, was from private investment. SOURCE: QUID

MILLIONS OF GITHUB AI PROJECTS

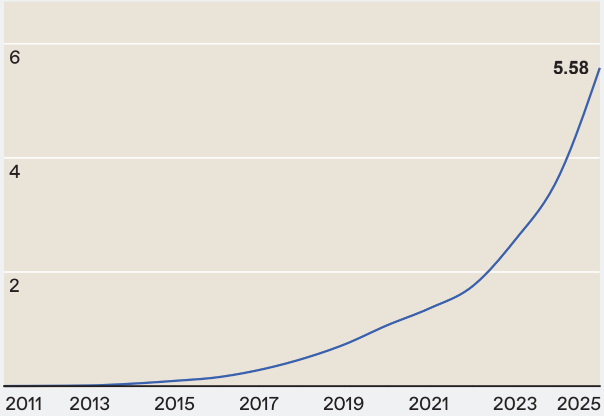

INDIVIDUALS ARE IN ON AI, TOO: The number of GitHub projects related to AI has soared since 2019, headlined by incredibly popular projects such as OpenClaw. Although some of the traffic is likely from AI agents, analysis suggests that the majority is still human activity. SOURCE: GITHUB

SEMICONDUCTORS

Wi-Fi Receiver Can Survive Inside a Nuclear Reactor ›

Wireless comms could unleash robots for decommissioning

BY KATHERINE BOURZAC

After the 2011 nuclear disaster at the Fukushima Daiichi plant, engineers began using robots to help assess the site. But due to the highly radioactive environment, those robots often relied on wired communication connections, which could get tangled or stuck.

Now, researchers in Japan have made a 2.4-gigahertz Wi-Fi receiver that's tough enough to work inside a nuclear reactor. They hope the receiver might become part of a wireless communications system for robotics used to decommission reactors.

Yasuto Narukiyo, a graduate student at the Institute of Science Tokyo, presented the wireless receiver in February at the IEEE International Solid-State Circuits Conference, in San Francisco. The receiver endured a total radiation dose of 500 kilograys, orders of magnitude higher than the doses typically tolerated by electronics in outer space (another high-radiation environment).

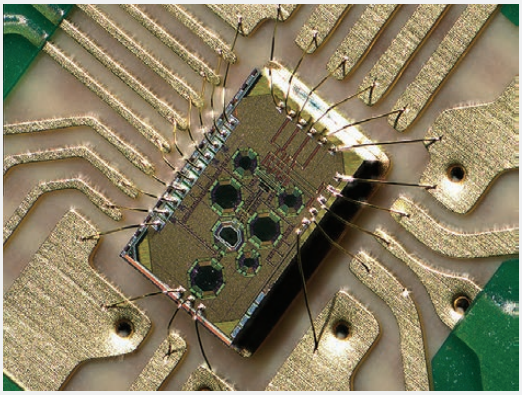

Researchers in Japan engineered a Wi-Fi receiver chip to survive in high-radiation environments, such as a nuclear reactor.

Even under less dramatic circumstances than Fukushima, nuclear plants don't last forever, and they need to be safely dismantled and decontaminated at the end of their lives. The process is lengthy and risks exposing people to radiation, which is why engineers hope robots can assist.

And the need is growing: According to a 2024 study, of 204 reactors that have been closed, only 11 plants with capacities over 100 megawatts have been fully decommissioned, and 200 more reactors will reach the end of their lifetimes in the next 20 years.

While electronics for space exploration are typically required to endure radiation doses of 100 to 300 grays over three years, a robot operating in a nuclear reactor needs to endure more than 500 kGy over the course of six months, says Narukiyo. For comparison, a robotic arm made by German industrial robotics manufacturer Kuka was able to withstand exposure to just under 165 Gy before failing.

To harden the 2.4-gigahertz Wi-Fi receiver against intense levels of radiation, the researchers, who also included Narukiyo's advisor Atsushi Shirane and Masaya Miyahara of Japan's High Energy Accelerator Research Organization, changed its mix of components, minimized the total number of transistors, and tinkered with the geometry of the transistors that were left.

The transistors, silicon MOSFETs (metal-oxide semiconductor field-effect transistors), contain an oxide layer that's particularly vulnerable to radiation damage. Blasts of gamma rays can trap positive charges in the oxide, degrading the device's performance and causing errors. So using fewer of them minimizes the problem. The researchers also made each transistor's gate longer and wider. The gate controls the flow of current—longer, wider gates perform better under exposure to radiation.

The group also considered the differences in how radiation affects PMOS transistors, MOSFETs in which current is carried primarily by positive charges, and NMOS, in which it is carried by electrons. PMOS transistors are more vulnerable to radiation damage because positive charge gets trapped in both the oxide and at the interface between the oxide and the rest of the semiconductor. These add up and shift the transistor toward the "off" state, says Narukiyo.

To compensate, the new receiver design minimizes the use of PMOS, replacing these transistors with other elements such as inductors that don't have an oxide layer. NMOS transistors are more resilient, says Narukiyo, because positive charges trapped in the oxide are to some extent canceled out by negative charges that get trapped at the interface.

Narukiyo and his team measured the performance of the receiver before exposure to radiation, and again after blasting it with total doses of 300 and 500 kGy. Before being irradiated, it showed comparable performance to typical Wi-Fi receivers. After reaching the highest radiation dose, the gain of the receiver had decreased by 1.4 decibels.

Narukiyo says the receiver is hardened enough, and now he hopes to improve its performance. He's also working on adding a transmitter, which would allow for two-way communications. This is more challenging due to the need to produce high levels of current to generate the Wi-Fi signal; a transmitter tested earlier was broken by a 300 kGy dose. The group is exploring using other semiconductors, such as diamond, to toughen the transmitter. ■

YASUTO NARUKIYO, SENA KATO, ET AL.

FLAPPER NIMBLE+

Flapper Nimble+ is a bioinspired flapping-wing drone that is safe around humans. It can take off and land vertically, hover, and fly in any direction. It is used as an entertainment robot in drone shows and also as a research platform.

CREATOR
Flapper Drones

YEAR
2021

COUNTRY
Czech Republic

CATEGORIES
Drones,
Entertainment,
Research

FEATURES
Agility,
full autonomy,
payload,
teleoperation,
wings

PHOTO: FLAPPER DRONES

To learn more about Flapper Nimble+ and explore over 270 real-world robots, with thousands of photos, videos, and interactive animations, scan the QR code or visit the IEEE Robots Guide at robotsguide.com

Hands On

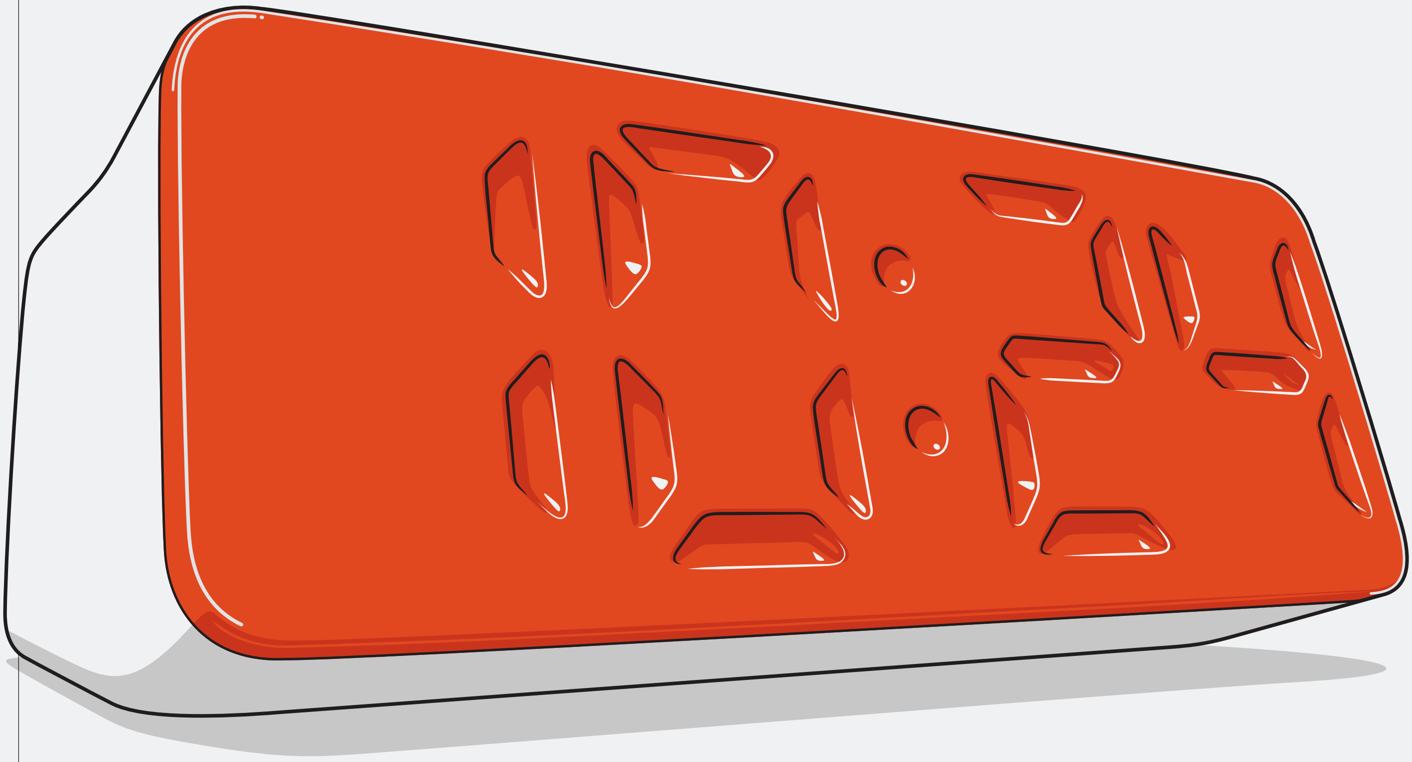

This clock's soft display relies on millifluidic logic and memory.

The Soft Clock >

This timepiece merges electronics and millifluidics

BY NILS JANßEN

Electrons are *great*. We use them to move vehicles, illuminate cities, and, of course, compute. But computation is not confined to the world of electronics. And shifting to alternative nonelectronic realms can unlock unique advantages: Photonic chips, for instance, process information with light while generating little heat. Another compelling alternative is fluidics, which uses pressurized gases or liquids to build logic circuits. Pioneered in the 1960s but sidelined by microchips, the field reemerged in the 1990s as "microfluidics." This approach aims to shrink laboratories onto a single chip by creating microscopic fluid channels with integrated micropneumatic control systems.

Today, there is a second fluidic revival, this time in the domain of soft robotics. Scaling microfluidic designs up to the millimeter-scale range (millifluidics) enables the higher flow rates necessary to drive robotic actuators. These robots exploit the nonlinear behaviors of soft materials to create lifelike motion and safer interactions, often utilizing pressurized air.

By building systems that "think" with the same air that powers them, we can drastically reduce the need for bulky electronic-to-pneumatic interfaces. This is the focus of my Soiboi Studio robotics lab. With millifluidic logic, I have steadily scaled the complexity of

Illustrations by James Provost

my designs. What began with a simple oscillator has most recently evolved into a clock featuring a soft, four-digit, seven-segment display.

Building on microfluidics research from the early 2000s and recent developments from the Grover Lab at the University of California, Riverside, I've created millifluidic devices using standard 3D printing and silicone casting. The basic architecture is simple: A flexible membrane is sandwiched between rigid layers embedded with networks of air channels.

Just as electronics rely on differing voltage potentials, these fluidic circuits operate on the pressure difference between atmospheric pressure (logical 0) and a near-vacuum at around −60 kilopascals of relative pressure (logical 1). Using negative pressure means the membrane is pulled into openings. This creates robust seals that allow me to replicate electronic building blocks.

While fluidic resistors are easily realized by adjusting the channel geometry, the heart of the system is a valve that mimics a metal-oxide-semiconductor field-effect transistor, or MOSFET. This vacuum "transistor" features a flow layer with two chambers (the source and drain) divided by a central valve seat and a control layer containing a cavity (the gate). A membrane runs between the control and flow layers and normally prevents airflow between the source and drain chambers. To switch the transistor on, a vacuum is applied to the gate chamber, sucking the membrane into the cavity and lifting it off the seat. This opens a path for airflow, equivalent to closing an electric circuit. By adding a small aperture to the membrane, I created a check valve—the fluidic equivalent of a diode. By combining transistors and resistive "pull-down" channels, I can build a full suite of logic gates.

The original microfluidic designs that inspired me were fabricated from etched glass and milled acrylic. Adapting them for a standard 3D printer required reengineering the logic elements and mastering two critical fabrication techniques.

First, I need airtight prints, yet printed plastic is notoriously porous. By printing at elevated temperatures, slow speeds, and slight overextrusion, I was able to fill microscopic gaps. When you're using transparent filament, there's a handy visual indicator: The more transparent the plastic appears, the lower its porosity.

Second, I used glass for my print bed. By printing the upper and lower chambers directly against this bed, I got the interface surface to become mirror smooth. This finish is essential for creating reliable, airtight seals. A 0.3-millimeter silicone membrane is placed between the layers and secured with screws.

A cast silicone membrane forms the face of the clock [top], while behind it sit 3D-printed millifluidic blocks [middle rows]. An Arduino Uno controls driver boards that operate solenoids, which are connected to valves that are attached to a vacuum pump [bottom row].

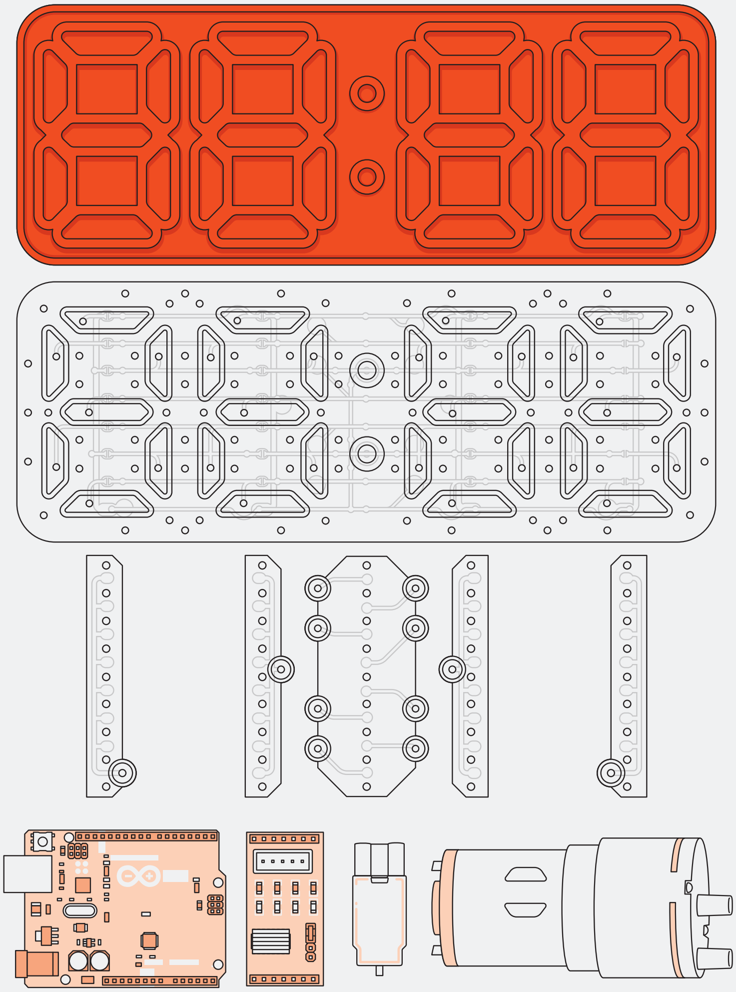

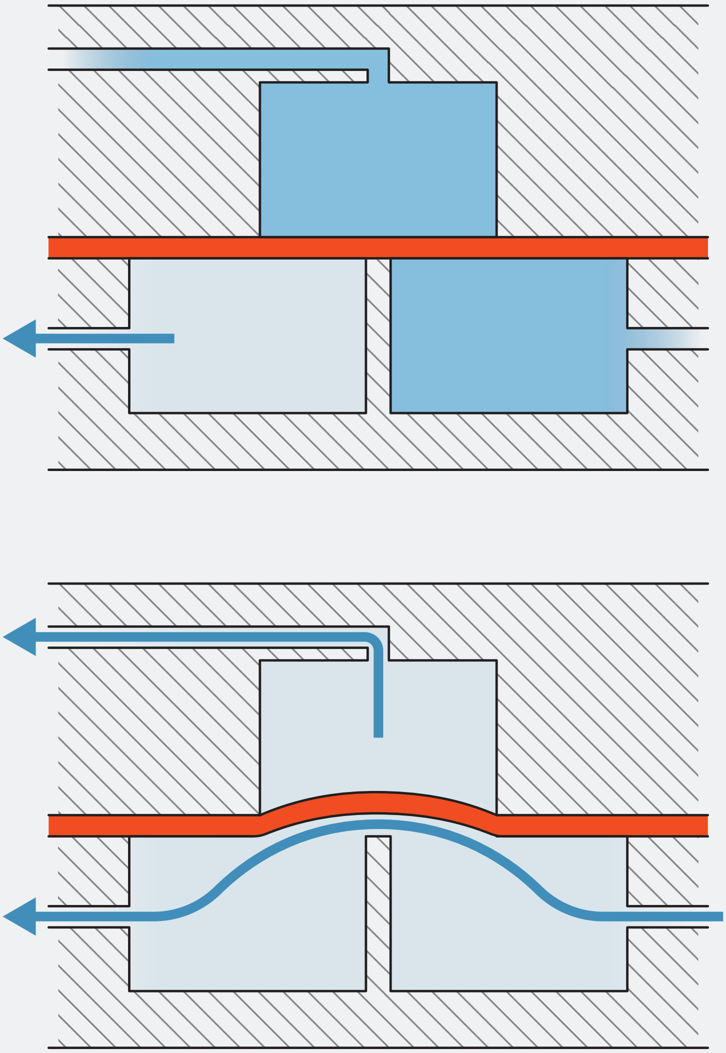

A pneumatic transistor is off when its upper control chamber is at atmospheric pressure [top]. When air is removed from the control chamber, it lifts a membrane, which allows air to move between the lower flow chambers and turns the transistor on [bottom].

The clockface is a cast silicone membrane. Each digit segment is formed by a small underlying cavity. When air is evacuated from this cavity, the membrane is sucked inward to create a concave hollow; when atmospheric pressure is restored, the silicone pops back flush with the surface. The result is a mesmerizing, organic motion.

The "brain" of the clock is an Arduino Uno, while the fluidics significantly reduce the hardware footprint. A four-digit, seven-segment display with two separator dots would require 29 solenoid valves to control directly. My clock needs just 11 valves.

To understand how it works, consider a standard electronic four-digit, seven-segment LED display. This also uses 11 pins to drive its digits. (In clockface displays, an additional pin is required to drive the separator dots.) Every digit is connected to a shared data bus with seven lines, one per segment. The four control lines select individual digits. Only one digit is illuminated at time, and strobing the digits at least 50 times per second creates the illusion that all four are simultaneously illuminated.

Such high-speed switching is not possible with air. Instead, I rely on memory. Each segment acts like a capacitor: By evacuating its cavity (logic 1), you "charge" the segment; by restoring atmospheric pressure (logic 0), you discharge it. Hence, each digit acts as an independent 7-bit memory. If the system is sufficiently airtight, the segments maintain their state for several seconds.

This display is more than a clock; it is a soft robot that happens to tell time.

Like the electronic display, the system utilizes a seven-line data bus. Each line connects to a solenoid valve that provides either vacuum or atmospheric pressure. To selectively address the individual digits, I placed a fluidic transistor between each segment and its data line. All the transistors' control inputs for a given digit are combined into one "write enable" line connected to its own solenoid valve. Activating this valve allows me to write data into the corresponding digit's memory.

The clock updates one digit per second, meaning a full cycle across the face takes 4 seconds. This cycle also drives the separator dots: A set of fluidic diodes connects the enable lines to the dots' cavities. Consequently, as each digit is addressed, the dots pulse automatically.

This display is more than a clock; it is a soft robot that happens to tell time. By offloading computation to the same air that powers movement, the clock approaches a new class of machines that are simpler, lighter, and more integrated. I'm now developing a guide for getting started with vacuum-powered logic and may release a refined version of this clock in the future. Watching the silicone skin morph serves as a fascinating reminder that not all logic needs silicon; sometimes, all you need is flexible silicone and a flow of air. ■

BY MAYSAM GHOVANLOO

Careers

Finding Success in Industry as a Chip Designer

Academe and industry require different approaches

Maysam Ghovanloo has spent his career designing application-specific ICs, first as an academic and now in industry.

SANDHYA MEEK/SILICON CREATIONS

I have been an application-specific IC (ASIC) designer for almost three decades. Over that time, I've moved through the full academic trajectory, from graduate student to full professor; later, I transitioned to industry after an unsuccessful stint at entrepreneurship. When I made the switch to the private sector in 2019, I began focusing on a critically important aspect of the electronic industry: silicon intellectual property.

As much as 80 percent of the physical area in today's most advanced chips is occupied by blocks that aren't made for specific products or even designed by the consumer-facing companies that built them. Instead, chipmakers draw heavily on established silicon IP from companies like Arm, Cadence, Rambus, Synopsys, and the company I work for, Silicon Creations.

Throughout my career, I've designed chips for very different purposes, including enabling the research program in my academic lab and expanding the IP portfolio of my company. When I joined Silicon Creations, I encountered a steep learning curve: I had no idea how much industry's approach to IC design differs from that of academia. Initially, it seemed that much of my two decades of academic research did not directly translate to my new role. I had to learn new skills and adopt a new mindset.

Today, demand for ASICs is rapidly growing, driven by the need for specialized chips in the automotive sector, AI applications, and more. By one market estimate, the ASIC market is expected to grow from US $23.4 billion to $38.8 billion by 2033, and the semiconductor industry as a whole is projected to hit $1 trillion by 2030. The industry needs more chip designers—but if you're coming from an academic background as I did, there are a few things you'll need to know.

The differences between industry and academe begin with a divergence in purpose. In academia, my primary objective was to generate new knowledge: to propose a novel circuit technique, validate an unconventional architecture, or explore the limits of performance in a given domain. A successful chip is one that demonstrates a concept. In industry, it is not nearly enough to prove that something can work. The goal is to ensure that it works reliably, repeatedly, and at scale. Success is measured not by novelty but by whether the silicon meets specifications, yields as expected in production, and supports a competitive product delivered on schedule.

This leads to a stark contrast in risk tolerance. Academic designs often deliberately push into unproven territory, where even partial success can yield valuable insight. In industry, however, we systematically minimize risk. The cost of failure makes first-time silicon success a central requirement—especially at advanced technology nodes, where the lithography masks used to transfer circuit designs onto silicon wafers alone can cost tens of millions of dollars. As a result, industry design flows are built around eliminating uncertainty through conservative margins, extensive validation, and careful reuse of proven solutions.

This paradigm has existed since the 1970s, when application-specific chip design was established. However, the gulf between academia and industry has expanded since the mid-2010s, when FinFET technology, a 3D architecture using vertical "fins" of silicon, was widely adopted in industry. System designs are also becoming increasingly modular with the advent of chiplets. This fundamentally altered the economics and complexity of ASIC development, with design costs rising by almost an order of magnitude. Initiatives like Taiwan Semiconductor Manufacturing Co.'s University FinFET Program and new government-funded chip-design hubs now let some well-resourced universities design for more advanced architectures, but the technology is still out of reach for many academics.

$1 trillion

MARKET VALUE OF THE SEMICONDUCTOR INDUSTRY EXPECTED BY 2030

Consider a startup developing an ASIC. Its engineering team may have deep expertise in a particular algorithm, sensor interface, or system architecture, the features that define its competitive advantage. But it is unlikely to possess world-class expertise in every supporting function. Developing each of these blocks internally would require significant time, capital, and specialized talent. Doing so could delay market entry beyond the startup's viability.

Even large semiconductor companies face similar constraints. Advanced-node development demands intense focus. Allocating a team to redesign a standard interface block that has already been implemented elsewhere may be difficult to justify when differentiation lies at the system level, such as an inference chip's ability to speed up neural network computations. The time it takes to move a new chip from conception to market and risk mitigation, not self-sufficiency, govern most decisions about in-house development versus outsourcing.

The economics of advanced IC manufacturing reinforce this reality. When the development cost of a leading-edge chip reaches hundreds of millions of dollars, minimizing risk becomes a central design imperative.

In this context, silicon IP emerged as a practical solution. Similar to how software developers rely on preexisting libraries rather than writing every function from scratch, ASIC designers license predesigned, preverified silicon blocks—such as processor cores, memory interfaces, and security engines—from highly specialized IP vendors. These blocks can then be integrated into larger, increasingly complex systems.

> **Academia explores the design space, asking what is possible, while industry exploits it, determining what is viable at scale.**

With the use of silicon IP, industry is able to widen the scope of its designs. Academic efforts tend to focus on block-level innovation: a new analog-to-digital converter architecture or an ultralow-noise amplifier, for instance. These designs typically abstract away many of the complexities of bringing a chip to market, such as packaging constraints, long-term reliability, and manufacturing yield.

In industry, the focus shifts to system-level integration. Modern systems on chips, or SoCs, incorporate dozens or even hundreds of functional blocks. Managing signal integrity, timing, firmware interaction, and system-level validation becomes as critical as the design of any individual block.

Verification philosophy also diverges sharply. In academia, the goal of verification is to demonstrate that the concept works under nominal conditions, which may not always reflect how it would perform in real applications. Even if only a fraction of fabricated chips from a multiproject wafer operates correctly, the design may still be considered a success if it validates the underlying idea.

At my academic lab for instance, we used to receive 40 chips from a TSMC prototyping service and started testing them in batches of five. If the first five or 10 chips proved functional, we had already collected more than enough data for a publication. If some of them failed, we weren't required to mention this when publishing the results.

In industry, verification is exhaustive, critical, and often dominates the development schedule. Failures are measured in parts per million, and even rare anomalies are carefully analyzed and documented to identify root causes and prevent recurrence. When I started at Silicon Creations, I was surprised by the level of detail and scrutiny designs face.

Differences in time horizons and economic constraints reinforce each of these contrasts. Academic projects operate on flexible timelines aligned with research and funding cycles. If I missed a deadline, I just had to wait for the next cycle. Industry projects are driven by fixed product schedules and market windows, frequently targeting costly leading-edge nodes to achieve competitive performance, power, and area efficiency. Missing a deadline can negate the value of an entire design and may have major financial consequences along the entire supply chain.

In essence, academia explores the design space, asking what is possible, while industry exploits it, determining what is viable at scale. Both are indispensable, but they operate under fundamentally different definitions of success. As ASIC complexity continues to grow, understanding both perspectives will be essential for the next generation of engineers navigating the evolving semiconductor landscape. ■

BY GWENDOLYN RAK

5 Questions

ELEANOR PARMENTER/CAMBRIDGE UNIVERSITY LIBRARY

Leontien Talboom

What it takes to preserve floppy disk data

Floppy disks are several decades old—many of the disks are degrading, and the data stored on them is at risk of being lost. In response, Leontien Talboom, a technical analyst at Cambridge University Libraries and Archives, led a project preserving floppy disks and their data called "Future Nostalgia," which concluded in January.

IEEE Spectrum spoke to Talboom about her work preserving data from Cambridge's collection of floppy disks and collecting knowledge about the disks themselves.

Why is it important to preserve floppy disk data?
Leontien Talboom: Two reasons. First, the physical media is starting to degrade. Floppy disks are made from plastic, but they've got a magnetic layer of iron oxide, and that's deteriorating. A lot of floppy disks are found in attics or garages, which means they also suffer from mold.

Second, a lot of people who developed floppy disks and systems that use floppy disks are starting to retire or pass away, which means that a lot of tacit knowledge is disappearing.

Leontien Talboom is a technical analyst at Cambridge University Libraries and Archives, where she transfers material from a wide range of storage media to make it accessible to archivists.

Whom did you go to for that tacit knowledge?
Talboom: I went to the retro computing community. Their work is more around preserving these machines to keep them running [than] the data that lives on the floppy disk. But they know their stuff about floppy disks.

For example, they know that in a lot of the older disks, the inside of the disk—the doughnut—gets stuck to the top. So if you flex the casing, the doughnut falls down again. If I hadn't known that, I would have assumed that those disks in our collection were broken or corrupt.

What is the most difficult part of working with floppy disks?
Talboom: Accessing the files can be quite challenging if we don't understand the file system. Within libraries and archives, we get a lot of material from machines that are not as well loved. Many of the personal computers that you had at home, such as the Amstrad or ZX Spectrum or BBC Micro, are very well documented. But a bunch of our material comes from business or research systems. They're not as nostalgic for people, so there's not as big a community preserving this type of material.

Do you have a favorite type of floppy disk?
Talboom: Five and a quarter. The weirder the system, the more frustrating and fun it is. I quite like doing that detective work.

The Amstrad disk has also really stolen my heart. The popularity of floppy disks is very geographically dependent. Our library, for example, has these Amstrad 3-inch disks. But if you go to the U.S., they're really uncommon. They weren't able to manufacture enough of these drives, and [3.5-inch disks] took over at a certain point. But they're really cute.

What's the best way to sustainably store data?
Talboom: The main thing is actively looking after it. A lot of the floppy disks we get in the library haven't been accessed for 20 or 30 years, which means that you need certain special hardware to actually read them, and then work with emulators or other tools to make these file formats accessible.

Now that we've done that work and transferred it, we can monitor it and make sure it's not suffering from anything like bit rot. We can also make decisions around migrating it to other file formats or working on specific file systems or unknown file formats in more detail. ■

Sheep return from pasture near the Bonorva wind farm, operated by EDF Renewables. The Florinas wind farm, commissioned in 2004, was one of the earliest wind farms built in Sardinia [right].

{SARDINIA'S ENERGY FUTURE HINGES ON ITS PAST}

Ancient and recent history keep the island stuck on coal

By EMILY WALTZ | *Photography by* LUIGI AVANTAGGIATO

“Why are you here?” Fabrizio Pilo, an electrical engineer, asks me as we sit in an outdoor café near his home in Cagliari, an ancient city on the island of Sardinia. It’s a fair question. I’m a journalist from the United States. I’d just stepped off my flight 2 hours prior and come straight to this meeting, suitcase still stowed in my rental car.

I’m here to see three intriguing new energy projects under development in Sardinia. I’d heard there’s strong public resistance to renewable energy, and I want to understand why that is. I tell Pilo, who is vice rector for innovation at the University of Cagliari, that I hope he’ll share some insights before I head out on a reporting trip across the island. (My answer seems to satisfy him, and he kindly gives me an hour of his time.)

This won’t be the first time that I’m asked to explain my presence on the island. I’d expected it, to some extent; I’m a foreign journalist poking around, after all.

What I didn’t expect was the depth of Sardinians’ distrust, not just of journalists, but of any outsider, particularly ones with authority. Over the last few years, developers of wind and solar projects, most of whom aren’t from here, have been absorbing the bulk of this smoldering, communal wariness.

In fact, the resistance is so widespread among Sardinians that over the course of two months in 2024, a grassroots petition to ban new wind and solar projects gathered over 210,000 certified signatures. That’s more than a quarter of Sardinia’s typical voter turnout and represents a cross-party consensus. People stood in long lines in public squares to sign. And it worked: Political leaders responded swiftly with an 18-month moratorium on renewable energy construction.

“I’ve never seen so much engagement for anything” in Sardinia, says Elisa Sotgiu, a literary sociologist at the University of Oxford, who was born and raised on the island. “Sardinia has a bunch of problems like enormous unemployment. There’s lots of emigration because there are no jobs. It’s one of the poorest areas in Europe. The area is just decaying,” she says. “And yet the thing people are demonstrating against is renewable energy.”

And the opposition continues: A network of mayors has mobilized for the cause. Thousands of people show up at organized protests. Activists vandalize grid equipment. Families are passing down these stories of resistance to their children as a point of pride. Local media outlets are egging it on, frequently publishing misinformation tinged with fearmongering.

These aren’t just NIMBY complaints—not in the pejorative sense, at least. The resistance, and the distrust underlying it, is rooted in the island’s complex history, both recent and ancient. It’s based on a past that the Sardinian people carry with them—a past that has seeded a deep sense of suspicion and vulnerability. Resistance, I learn, is part of what it means to be Sardinian.

“It is a very sad situation,” Pilo tells me. “There are a lot of economic reasons to do the [energy] transition.” It could attract new companies such as data centers, which would create new jobs, he argues. It could reduce Sardinia’s reliance on imported gas and fuel, making the island more independent. New economic activity on the island might help reverse its population decline, he adds.

And while what’s happening on Sardinia is unique, it also represents a larger trend: A growing number of communities around the world are opposing wind- and solar-farm construction, to the consternation of stakeholders. By 2025, nearly one-fourth of the counties in the United States had enacted some impediment to new utility-scale wind and solar energy—up from as few as 15 percent two years earlier, according to a *USA Today* analysis. In Africa, community pushback successfully canceled major projects such as the 60-megawatt Kinangop Wind Park in Kenya. In India, local pastoralists are challenging the 13-gigawatt Ladakh solar and wind project. And the European Union’s top-down push for renewable energy has created opposition in many communities.

Their reasons vary—land-use preferences, generational ethos, government resentment, property values, economic effects, aesthetics—but all of these struggles have this in common: The resisters are passionate and they are often successful in blocking development.

This is a looming problem for the energy transition. Unlike large, centralized coal and nuclear power plants, renewable energy is geographically spread out, so it touches far more communities. Sardinia offers one of the clearest cases of what can go wrong when renewable-energy developers and authorities fail to consider the complexities of the local situation on the ground.

ROUGHLY THE SIZE OF NEW HAMPSHIRE, Sardinia juts out of the Mediterranean Sea about 200 kilometers west of Italy’s mainland. Technically it’s part of Italy, but Sardinians are quick to point out their island’s autonomous status—a subtle way of saying, “We do things our way.” Its mountains seem to echo the sentiment. With the highest peaks running in a chain along the east side of the island, Sardinia resolutely turns its back to the mainland.

At first glance, the island looks like the kind of place that’s ripe for an energy transition. Its two coal plants are aging and are targeted to be shut down to meet climate commitments. It has no nuclear power, nor does it produce its own natural gas. Wind and sun, however, are abundant and could easily meet the energy needs of Sardinia’s sparse population of about 1.5 million.

{ NEWSPAPER *L’UNIONE SARDA* BOASTED IN 2024 THAT IT HAD ALREADY PUBLISHED MORE THAN A THOUSAND PAGES AGAINST RENEWABLE ENERGY SPECULATION SINCE 2022 AND PLANNED TO KEEP GOING. }

Concerned about the influx of solar and wind farms being built in Sardinia by outsiders, Roberto Pusceddu [left], under his pen name Erre Push, published a graphic novel that aimed to inspire young people to resist such impositions.

Fabrizio Pilo [above], vice rector of innovation at the University of Cagliari, has been working to help Sardinia transition to cleaner, more reliable energy.

But while the resources may be ready for a transition, the people emphatically are not. When I first arrive in Sardinia and take in its beauty, I assume that the impetus behind the fight against wind and solar farms boils down to how they look. Waves of silicon, metal, and concrete would spoil views of Sardinia's stunning beaches, rugged mountains, ancient pastures, and idyllic medieval villages, after all.

But the island's aesthetic—and the tourism industry that depends on it—are only part of the equation. The far stronger cultural forces at play are rooted in Sardinia's past. Over millennia, the island has endured successive invasions from outsiders seeking to exploit the land. These incursions, and Sardinians' rebellious responses to them, have become an integral part of the island's identity passed down through generations.

The invasions started with the relatively peaceful settlement of the Phoenicians in the 9th and 8th centuries B.C.E. Then came the Romans, the Byzantines, and the Iberians, who conquered with violence, looting, and enslavement. But legend has it that despite the might of these ancient conquerors, pockets of Sardinia sometimes managed to defend themselves. "Not even the Roman empire could conquer the shepherds of the highland regions," is the oft-repeated tale. Whether that's true or just an idealization is beside the point; such stories serve as an enormous source of pride and identity.

The island is "fiercely proud of its identity...especially in the center of Sardinia, which was the most resistant part," says Andrea Vargiu, a sociologist at the University of Sassari in Sardinia. "This long history of exploitation is still in our DNA, along with a proud sense of autonomy," he says.

Sardinia's unification, in the mid-1800s, with what would become the Kingdom of Italy is seen by many as an act of colonization. It didn't help that Italy then proceeded to exploit Sardinia's forests and other resources for the benefit of the mainland—a practice that continued through the 20th century, says Vargiu.

Sardinian bandits sometimes fought back with their own sense of justice, settling matters through raids, kidnappings, and violence. Their stories live on in Sardinian lore with an almost mythical quality, the brigands admired for their intractability.

Italy's use of the island for military purposes particularly irked locals. In a famous case in 1969, residents of the town of Orgosolo successfully thwarted the construction of a firing range on communal grazing land known as Pratobello. That name has since become synonymous with the defense of one's territory, and a rallying cry.

"Sardinia has always been a land of conquest," says Pasquale Mereu, mayor of Orgosolo, who spoke with *IEEE Spectrum* through an interpreter. "We believe that even today we are still

{ SARDINIA HAS ONE OF THE LARGEST CARBON FOOTPRINTS PER CAPITA IN ALL OF EUROPE. }

a colony of Italy, and I'm not ashamed to say it even though I represent an institution."

A longstanding mural on one of his village's walls reads: "You are in the territory of Orgosolo; here the people rule supreme and the government obeys."

DRIVING AROUND THE ISLAND and talking to people, I can feel the weight of Sardinia's history—and people's propensity for holding onto it. Elaborate heritage festivals occur nearly every autumn weekend in the island's interior. They're well attended, multigenerational affairs that aim to keep old traditions alive. In the medieval town of Belvì, men roast chestnuts—*marroni*—over an open fire in a frying pan the size of a swimming pool and then serve them to the crowd by shoveling them into troughs. They're delicious. In an adjacent amphitheater, the crowd sways along to costumed performers leading traditional dances.

Then there are the Bronze Age stone structures, called *nuraghi*, that are pretty much everywhere. Built before the violent conquests, these conical towers have come to symbolize a romanticized vision of the heyday of Sardinia's independence. More than 7,000 of them remain, ranging from unremarkable piles of rocks to complex towers, each one carefully documented on an interactive online map. I visit one of the more intact ones that's fenced off and requires an admission fee. As I take some video with my phone, an employee asks me who I am and what I'm doing and informs me I'll need to get permission from the government before posting anything online.

But in interviews with residents, I'm continually reminded of the darker side of Sardinia's past. People often bring up painful things that happened 50 or 500 years ago. A middle school science teacher named Giannina Serpi, and her husband, Roberto Moro, meet me at a café in the seaside town of Sant'Antioco. When I ask why people are so opposed to renewable energy, they (like many people I interviewed) point to the 1970s.

That decade brought a new kind of exploitation: not by empires or governments, but by technology companies. Petrochemical, aluminum, and other industrial companies from overseas built factories on the island, creating jobs and adjacent businesses. But after a few decades, economic and geopolitical factors led the companies to close the factories, sinking local economies and in some cases leaving behind toxic contamination.

In the northern city of Porto Torres, several petrochemical plants, a thermoelectric power plant, and an industrial harbor employed about 8,000 workers in the early 1970s. But the oil crises of that decade took their toll on jobs, and when environmental contamination became evident in the 1990s, employment plunged further. By 2010, most of the petrochemical plants had closed. Studies show that residents of Porto Torres during that time had curiously high rates of death from cancer, although there is no consensus on the cause.

Similarly, studies have found higher rates of lead in children in the Portovesme area in the southwest, about a 20-minute drive from where I sit with Serpi and Moro in Sant'Antioco. There, the U.S. aluminum producer Alcoa operated a smelter that employed about 500 people and supported an estimated 1,500 adjacent jobs. But the company shut down the smelter in 2012. Three years earlier, Russian aluminum manufacturer Rusal had idled its Eurallumina factory nearby.

The impacts of these events still feel fresh, Serpi explains through a digital translator. She says she teaches this history to her students but doesn't tell them how to feel about it. "I let them decide," she says.

Activists Maria Grazia Demontis [left] and Alberto Sala, photographed inside the archaeological monument Giants' Tomb of Pascarédda, have worked to stop the construction of wind farms by organizing protests and taking legal actions through their organization Gallura Coordination.

AGAINST THIS BACKDROP, renewable-energy developers in the early 2010s began sizing up Sardinia. They were drawn by the cheap land, low population, strong wind, and sun that shines an average of about 300 days a year. EF Solare Italia commissioned an 11-MW solar plant in 2010. Rome-based Enel Green Power began construction of a 90-MW wind farm in Portoscuso the following year.

Other developers followed, and they mostly came from elsewhere—mainland Italy, Europe, and later, China. The way many Sardinians saw it, the new plants didn't bring many long-lasting jobs. Most of the work ended after the design and installation phases, and profits went back to the companies' headquarters outside of Sar-

This rock hollowed out by erosion and walled up with stones was likely used by shepherds as a shelter near the historic Sardinian village of Tempio Pausania.

dinia, they argued. People called it "energy colonialism" and lauded landowners who refused to sell or lease their property to developers.

The uncle of Oxford's Sotgiu is one of those landowners. She says that a couple of years ago a solar company asked him if he would allow the installation of an array on his family farm in Logudoro in Sardinia's interior. "From that, he would have gotten something around €150,000 a year, which is more money than he's seen in his life," says Sotgiu. The money could have covered his three kids' college education, she says. "But he refused."

He had many reasons. For one, switching from sheep grazing to the more passive business of leasing land would have put the fate of his income in the hands of an outsider. "If you deprive a region of any sort of economy that is self-reliant, then it's really fragile," says Sotgiu. Her uncle didn't trust that the income would last, and worried he'd be left with a ruined farm, she says. Plus, his farm has been in the family for generations and one of his sons is interested in continuing the business. "So I understand his pride in saying, 'No, this is my farm, I don't care about the money,' " she says.

Despite that kind of grassroots resistance, development continued. In 2023, the Italian government authorized the construction of a 1-GW submarine power cable to connect Sardinia to Sicily and the Italian mainland. When completed, the bidirectional cable, called the Tyrrhenian Link, will increase electricity exchange between the regions, bolster grid reliability, and help grid operators efficiently use more renewable energy.

Sardinian activists, however, view the cable as a way to justify even more construction of wind and solar plants, and to export the island's energy for the benefit of non-Sardinians. The island already exports about 30 percent of its electricity, largely to Corsica and the Italian mainland via two existing submarine cables.

And then came the tipping point. In June 2024, in an effort to meet the European Union's 2030 renewable energy targets, Italy committed to building more than 80 GW over December 2020 levels of new wind and solar energy capacity. The national government divvied up the burden among its regions and told Sardinia to build its portion, 6.2 GW.

The move triggered an onslaught of requests from wind and solar developers wanting to build projects in Sardinia. The

Sardinia's grid is connected to nearby Corsica and the Italian mainland by two existing undersea power cables. Construction of a third subsea cable was completed this year.

Pink granite called Ghiandone Limbara was extracted from the Sinnada quarry in northern Sardinia from the late 1970s to 2011.
↓

MAP: ISTOCK

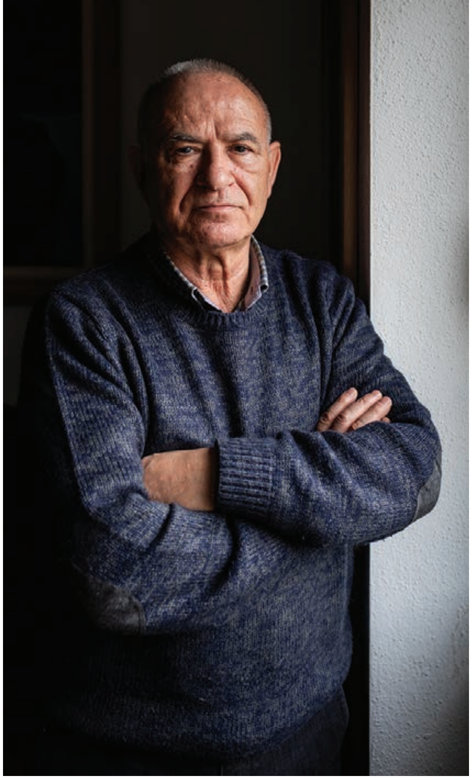

Pasquale Mereu, mayor of Orgosolo, helped organize the Pratobello 24 movement against renewable energy in Sardinia.

> { ITALY ORIGINALLY ORDERED SARDINIA'S TWO COAL PLANTS TO SHUT DOWN BY 2025 BUT LATER EXTENDED THE DEADLINE TO 2038. }

queue at one point topped 50 GW of grid-connection requests. That represented more than 700 solar and wind projects, many of which came from companies outside of Sardinia.

The southern newspaper *L'Unione Sarda* ran wild with the numbers. Almost daily, for months, it published stories about the "wind assault." The call-to-arms posts urged people to protest. "The Attack on the Landscape Does Not Stop; The Threat From Agrivoltaics Is Growing," read a July 2024 headline. Unsubstantiated articles tried to link wind and solar developers to organized crime.

"It was scaremongering," says Sotgiu. "It was a little dishonest, as I saw it, because they kept exaggerating and scaring people into thinking that we were going to be invaded." (Representatives of the newspaper declined to comment.)

The numbers did scare people. Lost was the fact that a grid-connection request is just the start of a multiyear process that involves permitting and legal review and often ends in withdrawn or downsized projects. Submitting a request is inexpensive, and developers often cast a wide net by entering lots of these queues globally to increase the odds of being accepted. In the end, only a fraction come to fruition. In other words, building all, or even most, of the requested 50 GW was never going to happen.

"I tried to explain this" to the public, says an industrial engineer at the University of Cagliari who asked to remain anonymous to avoid any detrimental impacts of speaking out. "I went to the regional television station. But it's difficult with technical information. And the newspaper communication is so bad, and its impact is so strong in the community, that it's very difficult to change people's minds," he says.

AND SO THE COLLECTIVE ANGST caused by powerful outsiders, industry, and the state united Sardinians into a singular cause. Faced with what felt like another attempted conquest, they did what their families and community had taught them to do: They resisted. Says Mayor Mereu: "This is what we are rebelling against: the idea that Sardinians are few and therefore must put up with everything."

In a nod to the 1969 resistance in Orgosolo, they dubbed the movement "Pratobello 2024." Activist groups, called "committees," organized protests, and created social media campaigns and videos. Thousands of people started showing up at planned demonstrations. A lawyer went on a hunger strike. Vandals unscrewed bolts on wind turbine blades and set fire to grid and construction equipment.

Italy's transmission system operator, Terna, had to switch to company cars without logos to avoid being targeted. Students studying the electricity system in a master's program sponsored by Terna were verbally attacked at an airport, according to a professor at their school who spoke with me about the violence.

Celebrities got involved. Italian actress and Bond Girl Caterina Murino met with Sardinia's president to ask her to reject wind farms. Murino posted on Instagram: "Nobody touch Sardinia!!!!" On Italian national TV, the jazz legend Paolo Fresu performed on trumpet while popular TV host Geppi Cucciari read an impassioned lament about the exploitation of the island.

Sardinian author Erre Push wrote a graphic novel titled *Fàula Birdi* about a protagonist who resisted an imposition from outsiders. He says he wrote it upon the request of the organization ReCommon, whose mission is to "challenge corporate and state power responsible for the plunder of territories." Push hopes the book will inspire more people to follow the protagonist's lead. "Renewables are another imposition like in the past—not to help Sardinians but to help external people like industry managers or founders of companies," he told me through an interpreter.

Mereu and a network of mayors drafted the petition that gathered so many signatures. The people had spoken. In response, Sardinian politicians passed a law that imposed an 18-month ban on construction of wind and solar projects within 7 km of a nuraghe or other archaeological site. It wasn't a total ban, but it might as well have been. "If you put a circle with a 7-km radius around each archaeological site, you cover all of Sardinia," says Emilio Ghiani, a power systems expert at the University of Cagliari. "In this way, it is impossible to find a place to install a new plant."

The move was like giving the Italian government—and the EU's clean energy targets—the middle finger. And it sent renewable-energy developers scrambling. One company building an agrivoltaic plant raced to bring construction to 30 percent completion, which the new law said was the threshold for being allowed to proceed. The company asked not to be named in this story to avoid trouble.

Furious, the government in Rome challenged the Sardinian regional law in Italy's Constitutional Court, and in January this year it prevailed. In its decision, the court rejected the law, saying that renewable-energy projects should be evaluated case by case.

Project development quickly resumed. So did the backlash. A headline in *L'Unione Sarda* declared: "Enough With Top-Down Decisions Without Consulting Communities."

WHERE THE ISLAND GOES FROM HERE is unclear. There's a willingness among a portion of the population to move forward with an energy transition. For example, some of Sardinia's largest cheese makers are powering their operations with renewable energy and installing systems to utilize waste heat for efficiency. But for the most part, the public isn't budging in its resistance. Researchers are trying to dispel inaccurate information, but regional newspapers seem bent on perpetuating fear.

Plus, there are technical issues to work out before a full-scale energy transition can be made. Sardinia's transmission system was built around the centralized generation of two coal plants; it wasn't made for the distributed generation of wind and solar plants. Renewables require a more dynamic grid, more energy storage, and a wider range of power sources to compensate for their intermittency. Engineers are working on it, but they've got a ways to go.

The new Tyrrhenian Link undersea power cable will help with those issues. By connecting Sardinia, Sicily, and the mainland, the cable creates more flexibility in the system. When wind or solar generation slows in Sardinia, for example, electricity from the mainland can fill in the gap, and vice versa. "It will increase the reliability of the system, and after it's installed, it will be possible to switch off the old generation plants that use coal," says Ghiani. In January, Terna finished laying the western section of the cable between Sardinia and Sicily, and in April it completed the eastern section between Sicily and Campania on the mainland. Doing so set a world record for power cable depth, at 2,150 meters below sea level, according to Terna.

The link is one of the most innovative high-voltage direct current (HVDC) projects in Europe. It can move up to a gigawatt of power and reverse that power flow nearly instantaneously. By using voltage source converter (VSC) technology, it can also help prevent power-flow problems by regulating frequency and smoothing out oscillations in the grid in real time. And it has black-start capability: In the event of a shutdown, it can help restore the grid without relying on an external electric network. These features are particularly helpful for an isolated network like Sardinia's.

Italy has created new incentives and regulations to build a market for grid-scale energy storage. Having plenty of storage is a key to scaling up renewables because it provides backup power when the wind isn't blowing or the sun isn't shining. To this end, Italy created MACSE, an auction that gives storage developers revenue certainty. Its name translates to mechanism for the procurement of electricity storage capacity. The first auction round, in September, successfully awarded 10 GWh.

Energy experts in Sardinia are also working with policymakers to change the rules around grid-connection requests. But these kinds of nerdy details don't grace most household conversations.

CONTINUED ON PAGE 50

Residents of the city of Orgosolo in 1969 famously stopped the construction of a military firing range on communal grazing land known as Pratobello. Its village walls are still covered in murals advocating social protest and antiauthoritarianism.

By **Olivia Hsu** & **Kalhan Koul**

Better Hardware Could Turn **Zeros** Into **AI Heroes**

Sparse computing enables leaner, faster AI

Illustration by Petra Péterffy

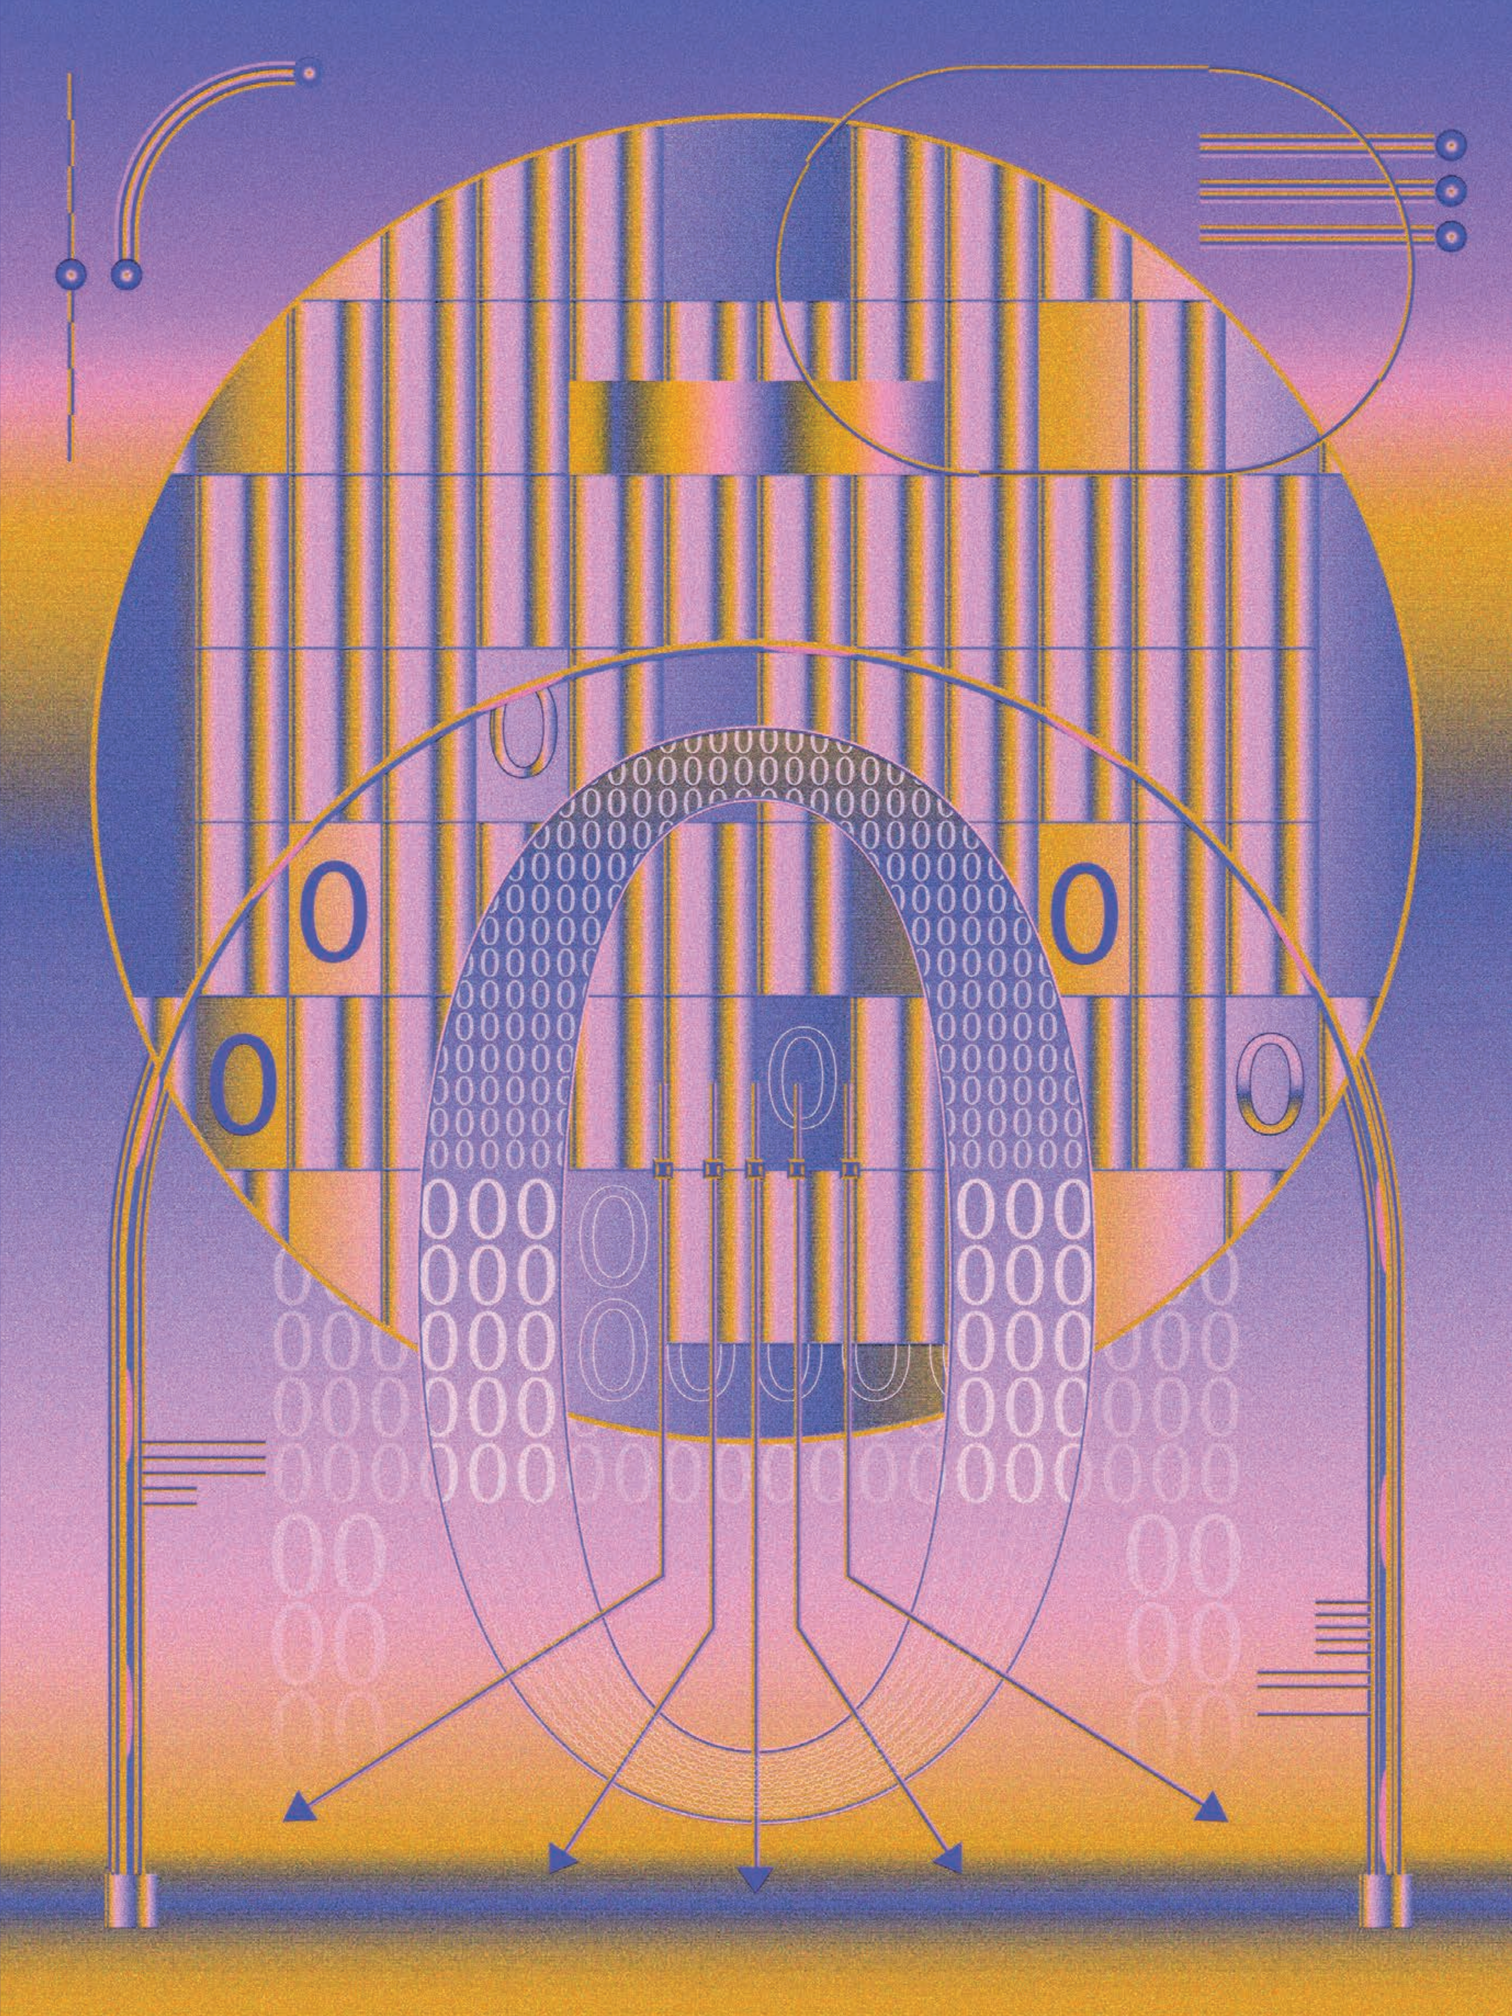

When it comes to AI models, size matters.

Even though some artificial-intelligence experts warn that scaling up large language models (LLMs) is hitting diminishing performance returns, companies are still coming out with ever larger AI tools. Meta's latest Llama release had a staggering 2 trillion parameters that define the model.

As models grow in size, their capabilities increase. But so do the energy demands and the time it takes to run the models, which increases their carbon footprint. To mitigate these issues, people have turned to smaller, less capable models and to using lower-precision numbers whenever possible for the model parameters.

But there is another path that may retain a staggeringly large model's high performance while reducing the time it takes to run and the size of the energy footprint. This approach involves befriending the zeros inside large AI models.

For many models, most of the parameters—the weights and activations—are actually zero, or so close to zero that they could be treated as such without losing accuracy. This quality is known as sparsity. Sparsity offers a significant opportunity for computational savings: Instead of wasting time and energy adding or multiplying zeros, these calculations could simply be skipped; rather than storing lots of zeros in memory, one need only store the nonzero parameters.

Unfortunately, today's popular hardware, like multicore CPUs and GPUs, do not naturally take full advantage of sparsity. To fully leverage sparsity, researchers and engineers need to rethink and re-architect each piece of the design stack, including the hardware, low-level firmware, and application software.

In our research group at Stanford University, we have developed the first (to our knowledge) piece of hardware that's capable of calculating all kinds of sparse and traditional workloads efficiently. The energy savings varied widely over the workloads, but on average our chip consumed one-seventieth the energy of a CPU, and performed the computation on average eight times as fast. To do this, we had to engineer the hardware, low-level firmware, and software from the ground up to take advantage of sparsity. We hope this is just the beginning of hardware and model development that will allow for more energy-efficient AI.

Neural networks, and the data that feeds into them, are represented as arrays of numbers. These arrays can be one-dimensional (vectors), two-dimensional (matrices), or more (tensors). A sparse vector, matrix, or tensor has mostly zero elements. The level of sparsity varies, but when zeros make up more than 50 percent of any type of array, it can stand to benefit from sparsity-specific computational methods. In contrast, an object that is not sparse—that is, it

SPARSE STORAGE AND COMPUTATION

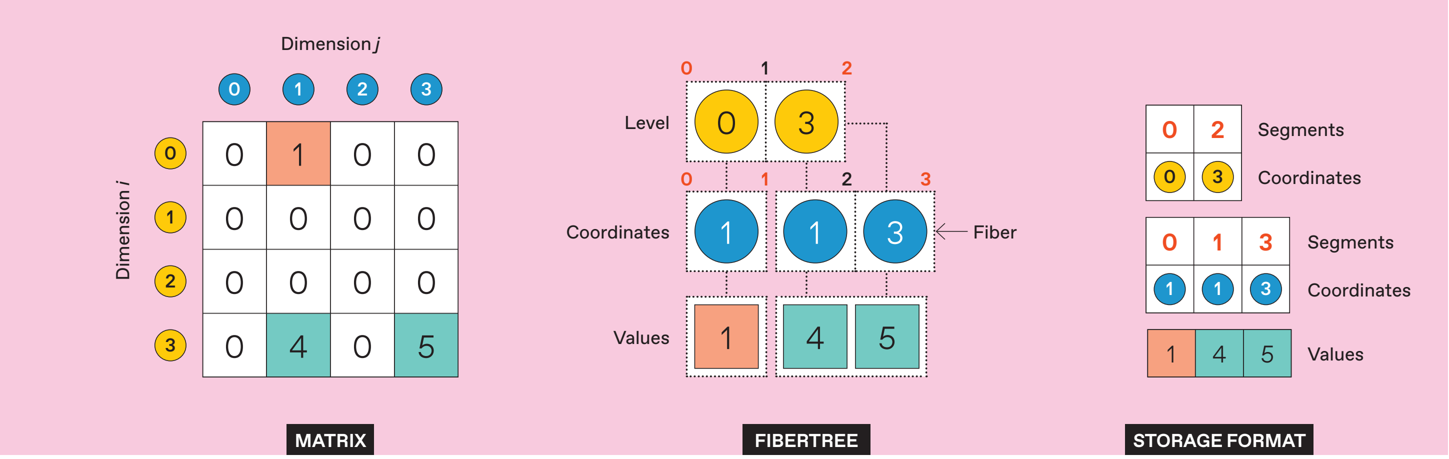

Normally, a four-by-four matrix takes up 16 spaces in memory, regardless of how many zero values there are. If the matrix is sparse, meaning a large fraction of the values are zero, the matrix is more effectively represented as a fibertree: a "fiber" of *i* coordinates representing rows that contain nonzero elements, connected to fibers of *j* coordinates representing columns with nonzero elements, finally connecting to the nonzero values themselves. To store a fibertree in computer memory, the "segments," or endpoints, of each fiber are saved alongside the coordinates and the values.

DENSE MATRIX-VECTOR MULTIPLICATION

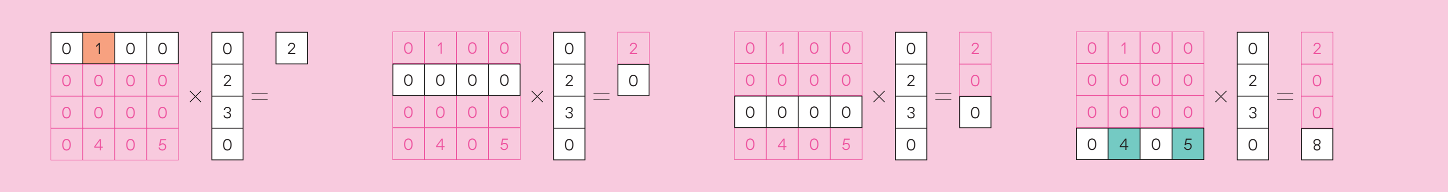

$$\begin{bmatrix}0&1&0&0\\0&0&0&0\\0&0&0&0\\0&4&0&5\end{bmatrix}\times\begin{bmatrix}0\\2\\3\\0\end{bmatrix}=\begin{bmatrix}2\end{bmatrix}\quad\begin{bmatrix}0&1&0&0\\0&0&0&0\\0&0&0&0\\0&4&0&5\end{bmatrix}\times\begin{bmatrix}0\\2\\3\\0\end{bmatrix}=\begin{bmatrix}2\\0\end{bmatrix}\quad\begin{bmatrix}0&1&0&0\\0&0&0&0\\0&0&0&0\\0&4&0&5\end{bmatrix}\times\begin{bmatrix}0\\2\\3\\0\end{bmatrix}=\begin{bmatrix}2\\0\\0\end{bmatrix}\quad\begin{bmatrix}0&1&0&0\\0&0&0&0\\0&0&0&0\\0&4&0&5\end{bmatrix}\times\begin{bmatrix}0\\2\\3\\0\end{bmatrix}=\begin{bmatrix}2\\0\\0\\8\end{bmatrix}$$

SPARSE MATRIX-VECTOR MULTIPLICATION

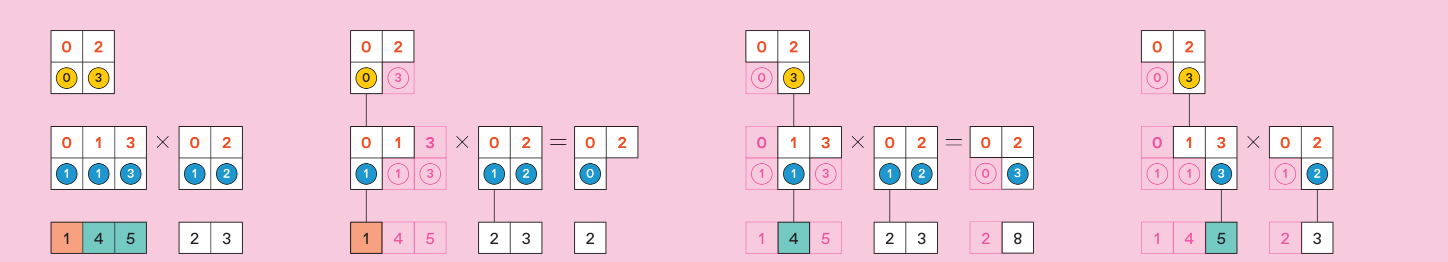

STEP 1: Look up the first nonzero values and find a match (both coordinates are equal to 1). Perform one multiplication.

STEP 2: Look up the next nonzero values and find a match (both coordinates are equal to 1). Perform one multiplication.

STEP 3: Look up next nonzero values and find no overlap. You're done!

Multiplying a vector by a matrix traditionally takes 16 multiplication steps and 16 addition steps. With a sparse number format, the computational cost depends on the number of overlapping nonzero values in the problem. Here, the whole computation is accomplished in three lookup steps and two multiplication steps.

DATA SOURCE: OLIVIA HSU & KALHAN KOUL

has few zeros compared with the total number of elements—is called dense.

Sparsity can be naturally present, or it can be induced. For example, a social-network graph will be naturally sparse. Imagine a graph where each node (point) represents a person, and each edge (a line segment connecting the points) represents a friendship. Since most people are not friends with one another, a matrix representing all possible edges will be mostly zeros. Other popular applications of AI, such as other forms of graph learning and recommendation models, contain naturally occurring sparsity as well.

Beyond naturally occurring sparsity, sparsity can also be induced within an AI model in several ways. Two years ago, a team at Cerebras showed that one can set up to 70 to 80 percent of parameters in an LLM to zero without losing any accuracy. Cerebras demonstrated these results specifically on Meta's open-source Llama 7B model, but the ideas extend to other LLM models like ChatGPT and Claude.

Sparse computation's efficiency stems from two fundamental properties: the ability to compress away zeros and the convenient mathematical properties of zeros. Both the algorithms used in sparse computation and the hardware dedicated to them leverage these two basic ideas.

First, sparse data can be compressed, making it more memory efficient to store "sparsely"—that is, in something called a sparse data type. Compression also makes it more energy efficient to move data when dealing with large amounts of it. This is best understood by an example. Take a four-by-four matrix with three nonzero elements. Traditionally, this matrix would be stored in memory as is, taking up 16 spaces. This matrix can also be compressed into a sparse data type, getting rid of the zeros and saving only the nonzero elements. In our example, this results in 13 memory spaces as opposed to 16 for the dense, uncompressed version. These savings in memory increase with increased sparsity and matrix size.

In addition to the actual data values, compressed data also requires metadata. The row and column locations of the nonzero elements also must be stored. This is usually thought of as a "fibertree": The row labels containing nonzero elements are listed and linked to the column labels of the nonzero elements, which are then linked to the values stored in those elements.

In memory, things get a bit more complicated still: The row and column labels for each nonzero value must be stored as well as the "segments" that indicate how many such labels to expect, so the metadata and data can be clearly delineated from one another.

In a dense, noncompressed matrix data type, values can be accessed either one at a time or in parallel, and their locations can be calculated directly with a simple equation. However, accessing values in sparse, compressed data requires looking up the coordinates of the row index and using that information to "indirectly" look up the coordinates of the column index before finally reaching the value. Depending on the actual locations of the sparse data values, these indirect lookups can be extremely random, making the computation data-dependent and requiring the allocation of memory lookups on the fly.

Second, two mathematical properties of zero let software and hardware skip a lot of computation. Multiplying any number by zero will result in a zero, so there's no need to actually do the multiplication. Adding zero to any number will always return that number, so there's no need to do the addition either.

In matrix-vector multiplication, one of the most common operations in AI workloads, all computations except those involving two nonzero elements can simply be skipped. Take, for example, the four-by-four matrix from the previous example and a vector of four numbers. In dense computation, each element of the vector must be multiplied by the corresponding element in each row and then added together to compute the final vector. In this case, that would take 16 multiplication operations and 16 additions (or four accumulations).

In sparse computation, only the nonzero elements of the vector need be considered. For each nonzero vector element, indirect lookup can be used to find any corresponding nonzero matrix element, and only those need to be multiplied and added. In the example shown here, only two multiplication steps will be performed, instead of 16.

Unfortunately, modern hardware is not well suited to accelerating sparse computation. For example, say we want to perform a matrix-vector multiplication. In the simplest case, in a single CPU core, each element in the vector would be multiplied sequentially and then written to memory. This is slow, because we can do only one multiplication at a time. So instead people use CPUs with vector support or GPUs. With this hardware, all elements would be multiplied in parallel, greatly speeding up the application. Now, imagine that both the matrix and vector contain extremely sparse data. The vectorized CPU and GPU would spend most of their efforts multiplying by zero, performing completely ineffectual computations.

Newer generations of GPUs are capable of taking some advantage of sparsity in their hardware, but only a particular kind, called structured sparsity. Structured sparsity assumes that, say, two out of every four adjacent parameters will be zero. However, some models benefit more from unstructured sparsity—the ability for any parameter (weight or activation) to be zero and compressed away, regardless of where it is and what it is adjacent to. GPUs can run unstructured sparse computation in software, for example, through the use of the cuSparse GPU library. However, the support for sparse computations is often limited, and the GPU hardware gets underutilized, wasting energy-intensive computations on overhead.

When doing sparse computations in software, modern CPUs may be a better alternative to GPU computation, because they are designed to be more flexible. Yet, sparse computations on the CPU are often bottlenecked by the indirect lookups used to find nonzero data. CPUs are designed to "prefetch" data based on what they expect they'll need from memory, but for randomly sparse data, that process often fails to pull in the right stuff from memory. When that happens, the CPU must waste cycles calling for the right data.

Apple was the first to speed up these indirect lookups by supporting a method called an array-of-pointers access pattern in the prefetcher of their A14 and M1 chips. Although innovations in prefetching make Apple CPUs more competitive for sparse computation, CPU architectures still have fundamental overheads that a dedicated sparse computing architecture would not, because they need to handle general-purpose computation.

Other companies have been developing hardware that accelerates sparse machine learning as well. These include Cerebras's Wafer Scale Engine and Meta's Training and Inference Accelerator (MTIA). The Wafer Scale Engine, and its corresponding sparse programming framework, have shown incredibly sparse results of up to 70 percent sparsity on LLMs. However, the company's hardware and software solutions support only weight sparsity, not activation sparsity, which is important for many applications. The second version of the MTIA claims a sevenfold sparse compute performance boost over the MTIA v1. However, the only publicly available information regarding sparsity support in the MTIA v2 is for matrix multiplication, not for vectors or tensors.

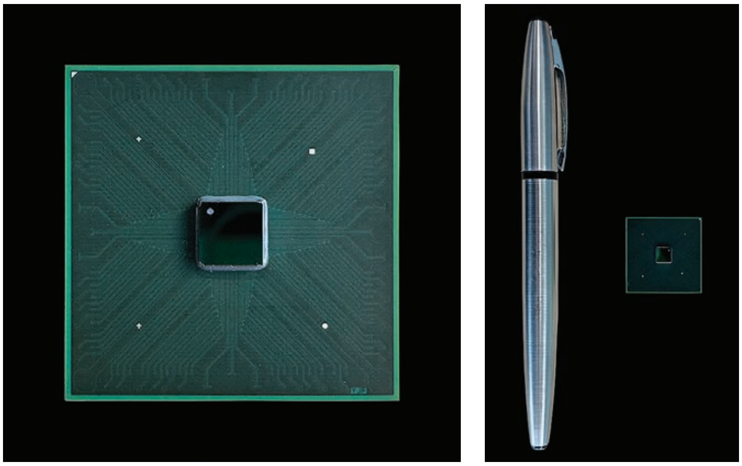
The Onyx chip, built on a coarse-grained reconfigurable array (CGRA), is the first (to our knowledge) to support both sparse and dense computations.

Although matrix multiplications take up the majority of computation time in most modern ML models, it's important to have sparsity support for other parts of the process. To avoid switching back and forth between sparse and dense data types, all of the operations should be sparse.

Instead of these halfway solutions, our team at Stanford has developed a hardware accelerator, Onyx, that can take advantage of sparsity from the ground up, whether it's structured or unstructured. Onyx is the first programmable accelerator to support both sparse and dense computation; it's capable of accelerating key operations in both domains.

To understand Onyx, it is useful to know

OLIVIA HSU

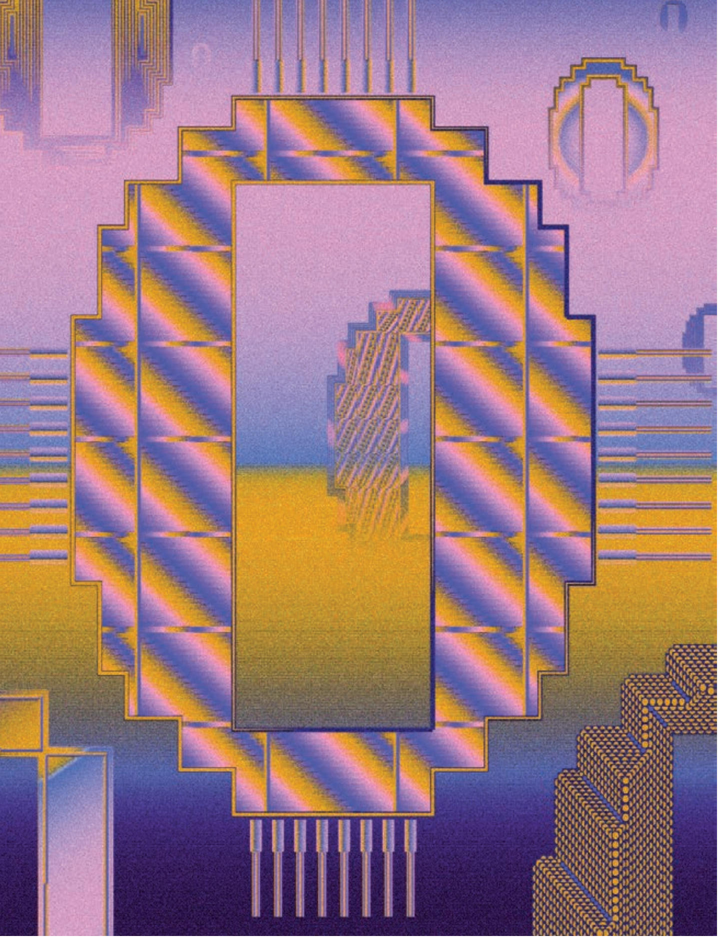

what a coarse-grained reconfigurable array (CGRA) is and how it compares with more familiar hardware, like CPUs and field-programmable gate arrays (FPGAs).

CPUs, CGRAs, and FPGAs represent a trade-off between efficiency and flexibility. Each individual logic unit of a CPU is designed for a specific function that it performs efficiently. On the other hand, since each individual bit of an FPGA is configurable, these arrays are extremely flexible, but very inefficient. The goal of CGRAs is to achieve the flexibility of FPGAs with the efficiency of CPUs.

CGRAs are composed of efficient and configurable units, typically memory and compute, that are specialized for a particular application domain. This is the key benefit of this type of array: Programmers can reconfigure the internals of a CGRA at a high level, making it more efficient than an FPGA but more flexible than a CPU.

Onyx is composed of flexible, programmable processing element (PE) tiles and memory (MEM) tiles. The memory tiles store compressed matrices and other data formats. The processing element tiles operate on compressed matrices, eliminating all unnecessary and ineffectual computation.

The Onyx compiler handles conversion from software instructions to CGRA configuration. First, the input expression—for instance, a sparse vector multiplication—is translated into a graph of abstract memory and compute nodes. In this example, there are memories for the input vectors and output vectors, a compute node for finding the intersection between nonzero elements, and a compute node for the multiplication. The compiler figures out how to map the abstract memory and compute nodes onto MEMs and PEs on the CGRA, and then how to route them together so they can transfer data between them. Finally, the compiler produces the instruction set needed to configure the CGRA for the desired purpose.

Since Onyx is programmable, engineers can map many different operations, such as vector-vector element multiplication, or the key tasks in AI, like matrix-vector or matrix-matrix multiplication, onto the accelerator.

We evaluated the efficiency gains of our hardware by looking at the product of energy used and the time it took to compute, called the energy-delay product (EDP). This metric captures the trade-off of speed and energy. Minimizing just energy would lead to very slow devices, and minimizing speed would lead to high-area, high-power devices.

Onyx achieves up to 565 times as much energy-delay product over CPUs (we used a 12-core Intel Xeon CPU) that utilize dedicated sparse libraries. Onyx can also be configured to accelerate regular, dense applications, similar to the way a GPU or TPU would. If the computation is sparse, Onyx is configured to use sparse primitives, and if the computation is dense, Onyx is reconfigured to take advantage of parallelism, similar to how GPUs function. This architecture is a step toward a single system that can accelerate both sparse and dense computations on the same silicon.

Just as important, Onyx enables new algorithmic thinking. Sparse acceleration hardware will not only make AI more performance- and energy efficient but also enable researchers and engineers to explore new algorithms that have the potential to dramatically improve AI.

Our team is already working on next-generation chips built off of Onyx. Beyond matrix multiplication operations, machine learning models perform other types of math, like nonlinear layers, normalization, the softmax function, and more. We are adding support for the full range of computations on our next-gen accelerator and within the compiler. Since sparse machine learning models may have both sparse and dense layers, we are also working on integrating the dense and sparse accelerator architecture more efficiently on the chip, allowing for fast transformation between the different data types. We're also looking at ways to manage memory constraints by breaking up the sparse data more effectively so we can run computations on several sparse accelerator chips.

We are also working on systems that can predict the performance of accelerators such as ours, which will help in designing better hardware for sparse AI. Longer term, we're interested in seeing whether high degrees of sparsity throughout AI computation will catch on with more model types, and whether sparse accelerators become adopted at a larger scale.

Building the hardware to optimally take advantage of zeros is just the beginning. With this hardware in hand, AI researchers and engineers will have the opportunity to explore new models and algorithms that leverage sparsity in novel and creative ways. We see this as a crucial research area for managing the ever-increasing runtime, costs, and environmental impact of AI. ■

Illustration by Petra Péterffy

Your Virtual Twin Could One Day Save Your Life

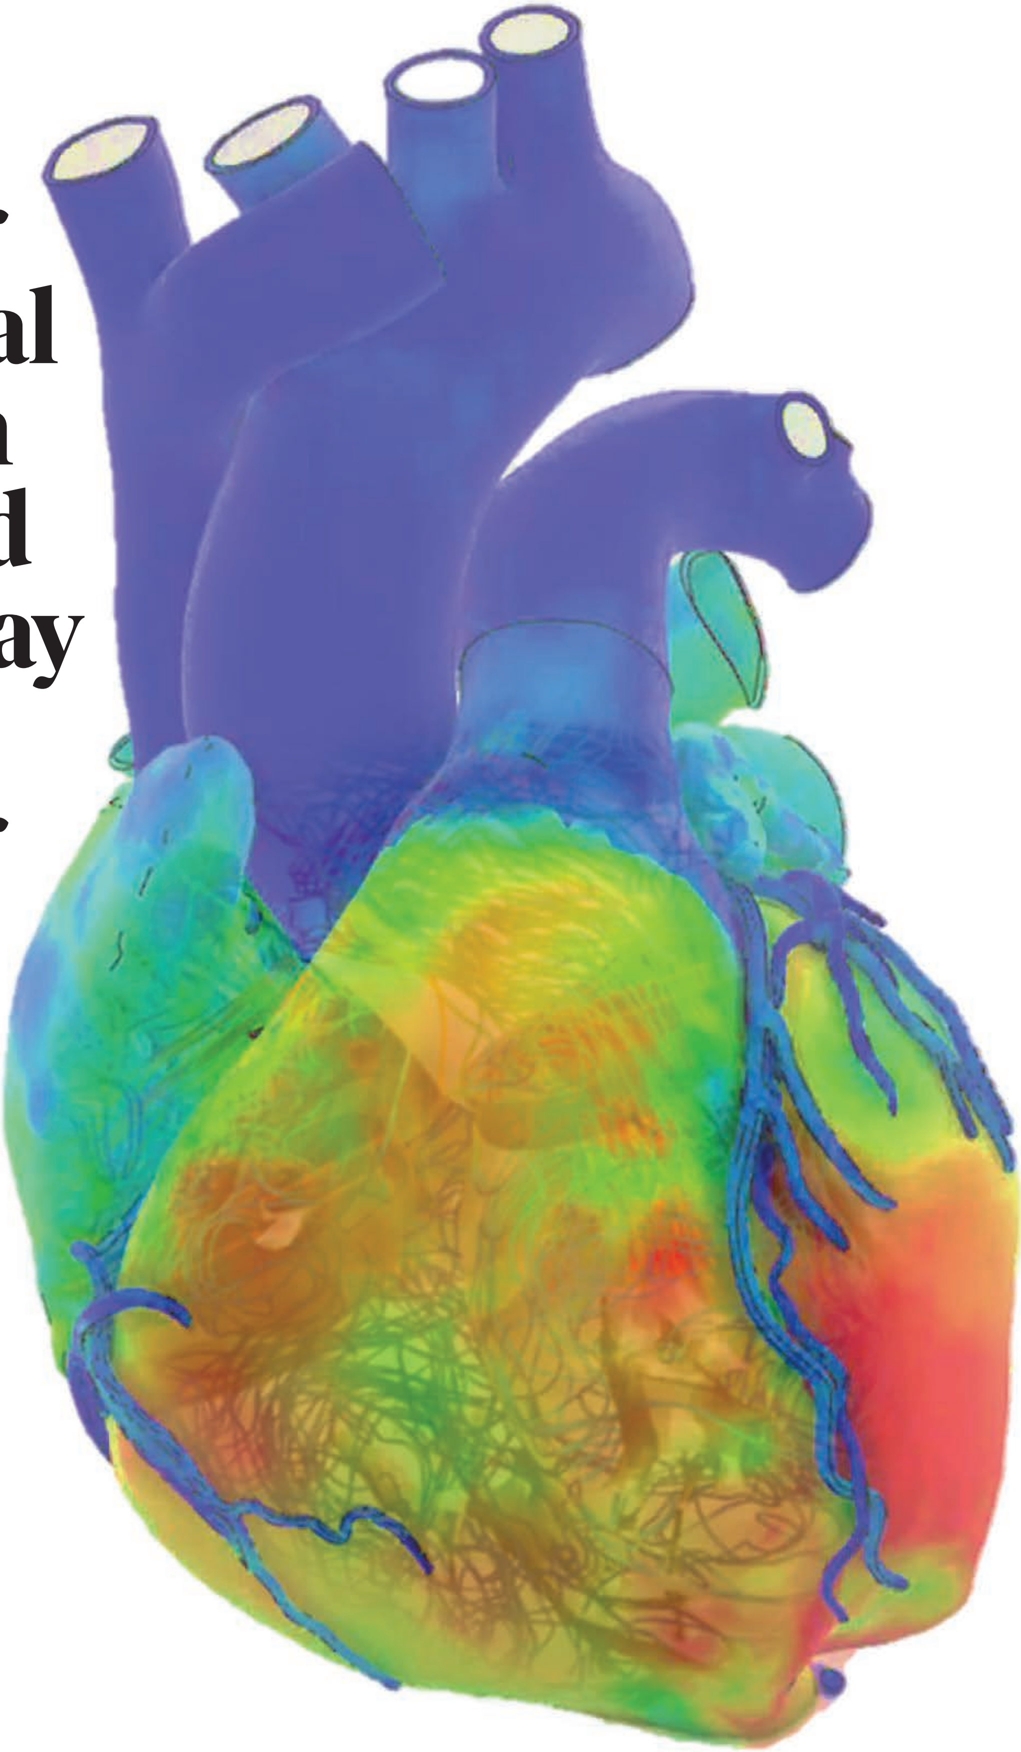

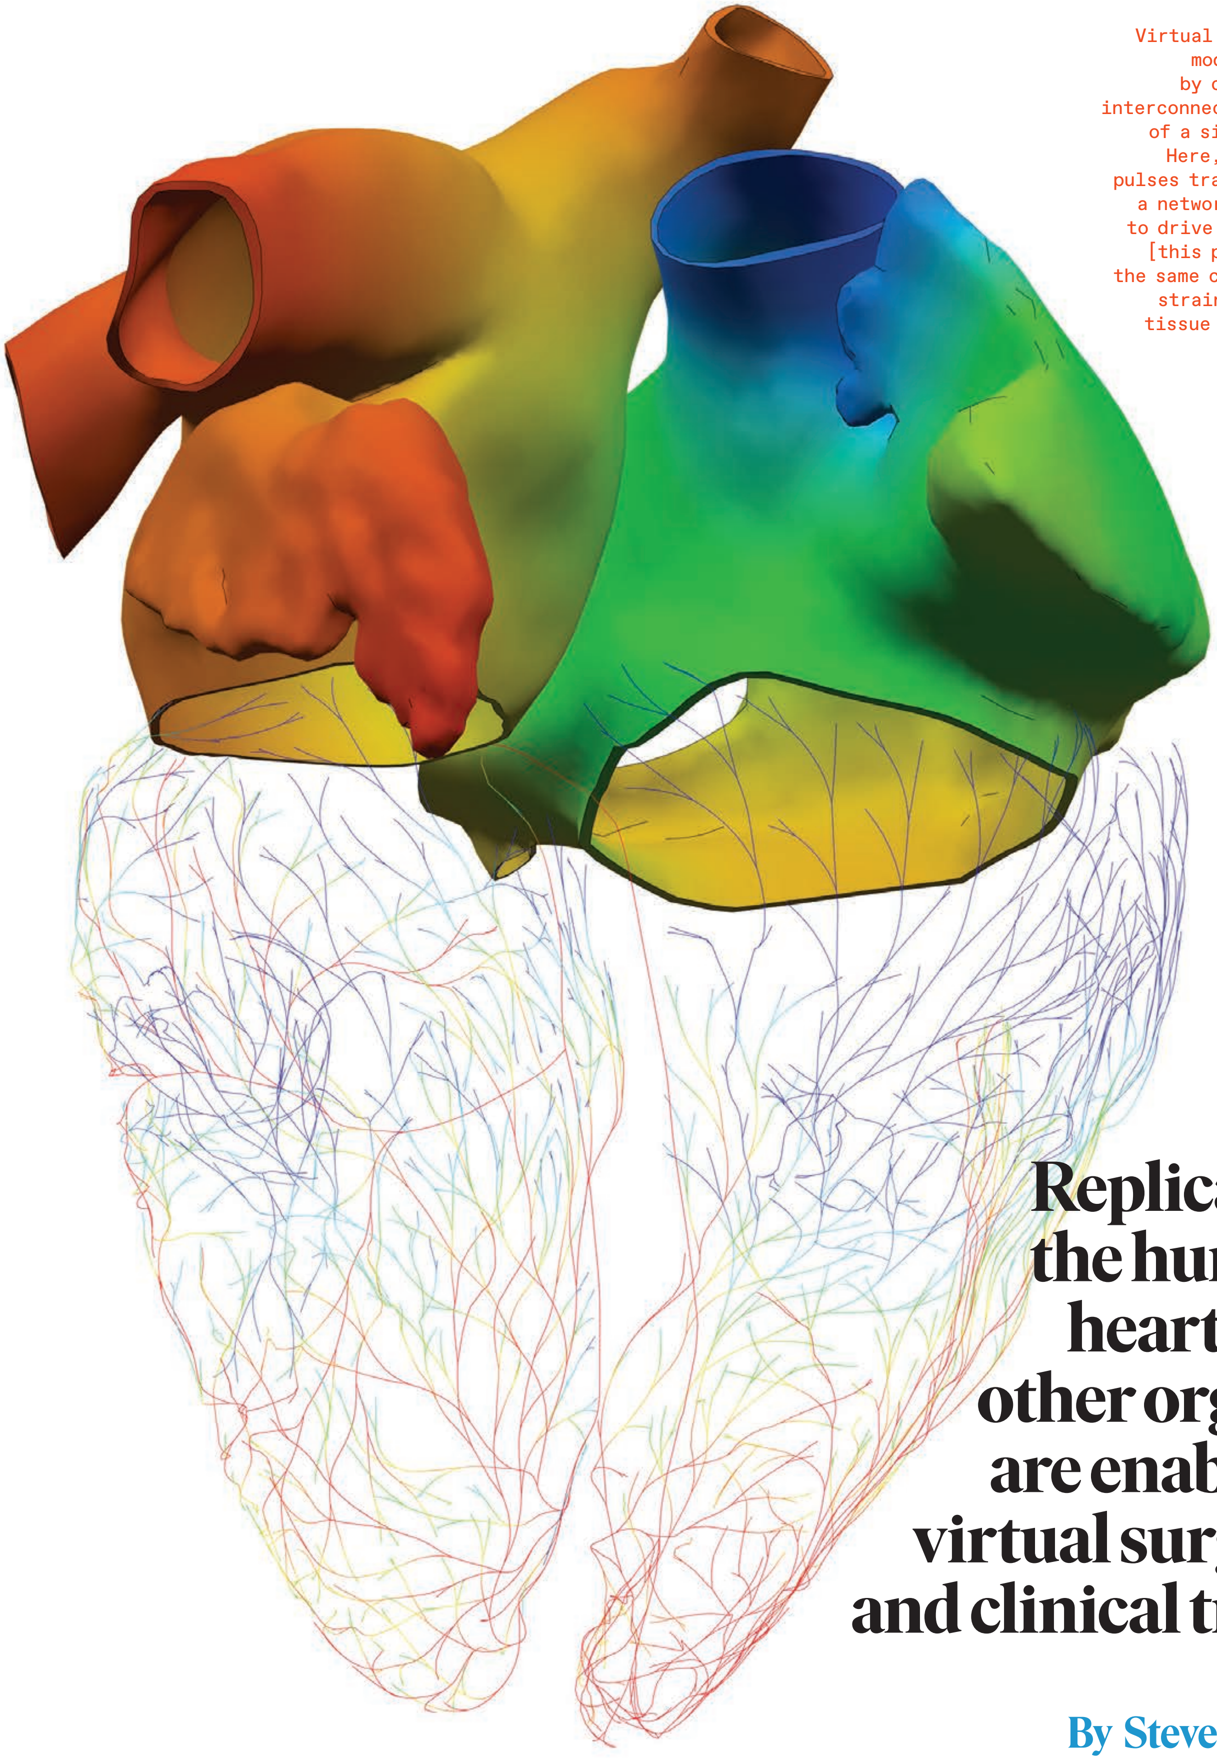

Virtual twins bring models to life by coupling the interconnected systems of a single organ. Here, electrical pulses travel through a network of fibers to drive a heartbeat [this page], while the same contractions strain the muscle tissue [opposite].

Replicas of the human heart and other organs are enabling virtual surgery and clinical trials

By Steve Levine

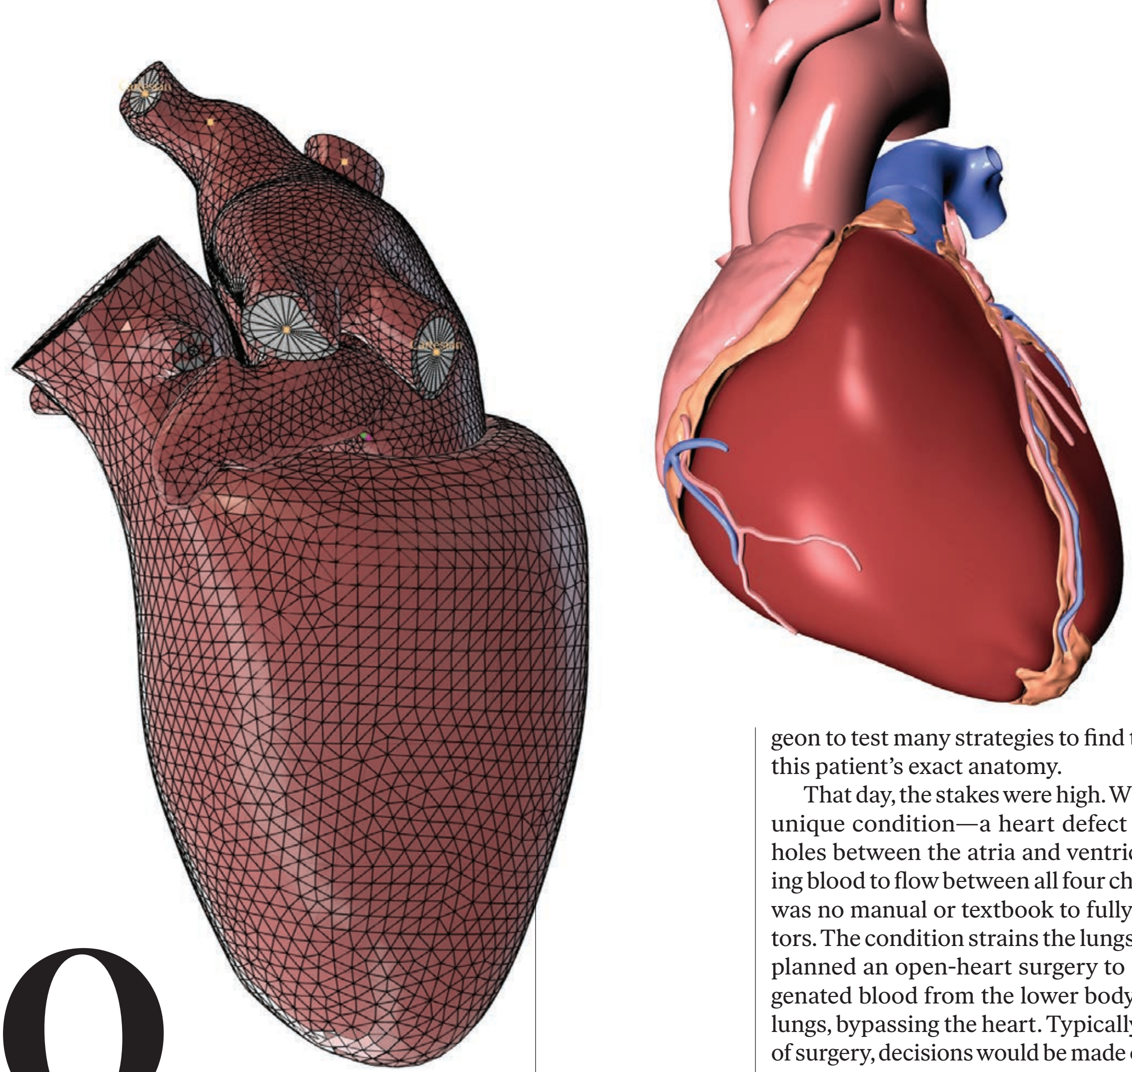

Each part of the heart, such as the left ventricle [left], is superimposed with a detailed digital mesh to re-create its physiology. These pieces come together to form an anatomically accurate rendering of the whole organ [right].

ALL IMAGES: DASSAULT SYSTÈMES

One morning in May 2019, a cardiac surgeon stepped into the operating room at Boston Children's Hospital more prepared than ever before to perform a high-risk procedure to rebuild a child's heart. The surgeon was experienced, but he had an additional advantage: He had already performed the procedure on this child dozens of times—virtually. He knew exactly what to do before the first cut was made. Even more important, he knew which strategies would provide the best possible outcome for the child whose life was in his hands.

How was this possible? Over the prior weeks, the hospital's surgical and cardio-engineering teams had come together to build a fully functioning model of the child's heart and surrounding vascular system from MRI and CT scans. They began by carefully converting the medical imaging into a 3D model, then used physics to bring the 3D heart to life, creating a dynamic digital replica of the patient's physiology. The mock-up reproduced this particular heart's unique behavior, including details of blood flow, pressure differentials, and muscle-tissue stresses.

This type of model, known as a virtual twin, can do more than identify medical problems—it can provide detailed diagnostic insights. In Boston, the team used the model to predict how the child's heart would respond to any cut or stitch, allowing the surgeon to test many strategies to find the best one for this patient's exact anatomy.

That day, the stakes were high. With the patient's unique condition—a heart defect in which large holes between the atria and ventricles were causing blood to flow between all four chambers—there was no manual or textbook to fully guide the doctors. The condition strains the lungs, so the doctors planned an open-heart surgery to reroute deoxygenated blood from the lower body directly to the lungs, bypassing the heart. Typically with this kind of surgery, decisions would be made on the fly, under demanding conditions, and with high uncertainty. But in this case, the plan had been tested in advance, and the entire team had rehearsed it before the first incision. The surgery was a complete success.

Such procedures have become routine at the Boston hospital. Since that first patient, nearly 2,000 procedures have been guided by virtual-twin modeling. This is the power of the technology behind the Living Heart Project, which I launched in 2014, five years before that first procedure. The project started as an exploratory initiative to see if modeling the human heart was possible. Now with more than 150 member organizations across 28 countries, the project includes dozens of multidisciplinary teams that regularly use multiscale virtual twins of the heart and other vital organs.

This technology is reshaping how we understand and treat the human body. To reach this transformative moment, we had to solve a fundamental challenge: building a digital heart accurate enough—and trustworthy enough—to guide real clinical decisions.

Now entering its second decade, the Living Heart Project was born in part from a personal conviction. For many years, I had watched helplessly as my daughter Jesse faced endless diagnostic uncertainty due to a rare congenital heart condition in which the position of the ventricles is reversed, threatening her life as she grew. As an engineer, I understood that the heart was an array of pumping chambers controlled by an electrical signal and its blood flow carefully regulated by valves. Yet I struggled to grasp the unique structure and behavior of my daughter's heart well enough to contribute meaningfully to her care. Her specialists knew the bleak forecast children like her faced if left untreated, but because every heart with her condition is anatomically unique, they had little more than their best guesses to guide their decisions about what to do and when to do it. With each specialist, a new guess.

Then my engineering curiosity sparked a question that has guided my career ever since: Why can't we simulate the human body the way we simulate a car or a plane?

I had spent my career developing powerful computational tools to help engineers build digital models of mechanical systems, using models that ranged from the interactions of individual atoms to the components of entire vehicles. What most of these models had in common was the use of physics to predict behavior and optimize performance. But in medicine today, those same physics-based approaches rarely inform decision-making. In most clinical settings, treatment decisions still hinge on judgments drawn from static 2D images, statistical guidelines, and retrospective studies.

This was not always the case. Historically, physics was central to medicine. Early doctors were, in a sense, applied physicists. They understood the heart as a pump, the lungs as bellows, and the body as a dynamic system. To be a physician meant you were a master of physics as it applied to the human body.

As medicine matured, biology and chemistry grew to dominate the field, and the knowledge of physics got left behind. But for patients like my daughter, that child in Boston, and millions like them, outcomes are governed by mechanics. No pill or ointment—no chemistry-based solution—would help, only physics. While I did not realize it at the time, virtual twins can reunite modern physicians with their roots, using engineering principles, simulation science, and artificial intelligence.

The LHP concept was simple: Could we combine what hundreds of experts across many specialties knew about the human heart to build a digital twin accurate enough to be trusted, flexible enough to personalize, and predictive enough to guide clinical care?

We invited researchers, clinicians, device and drug companies, and government regulators to share their data, tools, and knowledge toward a common goal that would lift the entire field of medicine. The Living Heart Project launched with a dozen or so institutions on board. Within a year, we had created the first fully functional virtual twin of the human heart.

The Living Heart was not an anatomical rendering, tuned to simply replicate what we observed. It was a first-principles model, coupling the network of fibers in the heart's electrical system, the biological battery that keeps us alive, with the heart's mechanical response, the muscle contractions that we know as the heartbeat.

Academic researchers had long explored computational models of the heart, but those projects were typically limited by the technology they had access to. Our version was built on industrial-grade simulation software from Dassault Systèmes, a company best known for modeling tools used in aerospace and automotive engineering, where I was working to develop the engineering simulation division. This platform gave teams the tools to personalize an individual heart model using the patient's MRI and CT data, blood-pressure readings, and echocardiogram measurements, directly linking scans to simulations.

Surgeons then began using the Living Heart to model procedures. Device makers used it to design and test implants. Pharmaceutical companies used it to evaluate drug effects such as toxicity. Hundreds of publications have emerged from the project, and because they all share the same foundation, the findings can be reproduced, reused, and built upon. With each application, the research community's understanding of the heart snowballed.

Early on, we also addressed an essential requirement for these innovations to make it to patients: regulatory acceptance. Within the project's first year, the U.S Food and Drug Administration agreed to join the project as an observer. Over the next several years, methods for using virtual-heart

At a visualization center in Boston, virtual reality imagery helps the mother of a young girl with a complex heart defect understand the inner workings of her child's heart.

models as scientific evidence began to take shape within regulatory research programs. In 2019, we formalized a second five-year collaboration with the FDA's Center for Devices and Radiological Health with a specific goal.

That goal was to use the heart model to create a virtual patient population and re-create a pivotal trial of a previously approved device for repairing the heart's mitral valve. This helped our team learn how to create such a population, and let the FDA experiment with evaluating virtual evidence as a replacement for evidence from flesh-and-blood patients. In August 2024, we published the results, creating the first FDA-led guidelines for in silico clinical trials and establishing a new paradigm for streamlining and reducing risk across the entire clinical-trial process.

In 10 years, we went from a concept that many people doubted could be achieved to regulatory reality. But building the heart was only the beginning. Now we've expanded the project to develop twins of other organs, including the lungs, liver, brain, eyes, and gut. Each corresponds to a different medical domain, which has its own community, data types, and clinical use cases. Working independently, these teams are progressing toward a breakthrough in our understanding of the human body: a multiscale, modular twin platform where each organ twin could plug into a unified virtual human.

A cardiac digital twin starts with medical imaging, typically an MRI, a CT, or both. The slices are reconstructed into the 3D geometry of the heart and connected vessels. The geometry of the whole organ must then be segmented into its constituent parts, so each substructure—atria, ventricles, valves, and so on—can be assigned their unique properties.

At this point, the object is converted to a functional, computational model that can represent how the various cardiac tissues deform under load—the mechanics. The complete digital twin model becomes "living" when we integrate the electrical fiber network that drives mechanical contractions in the muscle tissue.

To simulate circulation, the twin adds computational models of hemodynamics, the physics of blood flow and pressure. The model is constrained by boundary conditions of blood flow, valve behavior, and vascular resistance set to closely match human physiology. This lets the model predict blood flow patterns, pressure differentials, and tissue stresses.

Finally, the model is personalized and calibrated using available patient data, such as how much the volume of the heart chambers changes during the cardiac cycle, pressure measurements, and the timing of electrical pulses. This means the twin reflects not only the patient's anatomy but how their specific heart functions.

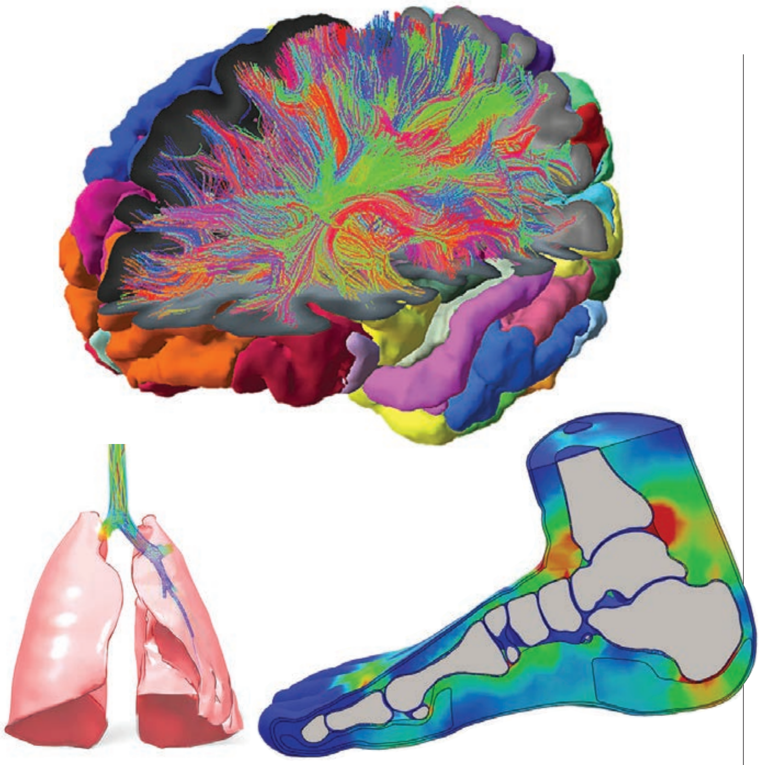

The Living Heart Project has begun modeling organs throughout the body. The 3D brain reconstruction [top] shows major pathways in the brain's white matter connecting color-coded regions of the brain. The cross section of a patient's foot [right] shows points of strain in the soft tissue when bearing weight. And the lung virtual twin [left] combines the organ's geometry with a physics-based simulation of air flowing down the trachea and into the bronchi.

When the FDA in silico clinical trial initiative launched in 2019, the project's focus shifted from these handcrafted virtual twins of specific patients to cohorts large enough to stand in for entire trial populations. That scale is feasible today only because virtual twins have converged with generative AI. Modeling thousands of patients' responses to a treatment or projecting years of disease progression is prohibitively slow with conventional digital-twin simulations. Generative AI removes that bottleneck.

AI boosts the capability of virtual twins in two complementary ways. First, machine learning algorithms are unrivaled at integrating the patchwork of imaging, sensor, and clinical records needed to build a high-fidelity twin. The algorithms rapidly search thousands of model permutations, benchmark each against patient data, and converge on the most accurate representation. Workflows that once required months of manual tuning can now be completed in days, making it realistic to spin up population-scale cohorts.

Second, enriching AI models' training sets with data from validated virtual patients grounds the AI simulations in physics. By contrast, many conventional AI predictions for patient trajectories rely on statistical modeling trained on retrospective datasets. Such models can drift beyond physiological reality, but virtual twins anchor predictions in the laws of hemodynamics, electrophysiology, and tissue mechanics. This added rigor is indispensable for both research and clinical care—especially in areas where real-world data are scarce, whether because a disease is rare or because certain patient populations, such as children, are underrepresented in existing datasets.

On the research side, the In Silico Clinical Trial Project we completed in 2024 opened a new world for medical innovations. A conventional clinical trial may take a decade, and 90 percent of new drug treatments fail in the process. Virtual twins, combined with AI methods, let researchers design and test treatments quickly. With a small library of virtual twins, AI models can rapidly create expansive virtual patient cohorts to cover any subset of the general population. As clinical data becomes available, it can be added into the training set to increase reliability and enable better predictions.

Virtual twin cohorts can represent a realistic population by building individual "virtual patients" that vary by age, gender, race, weight, disease state, comorbidities, and lifestyle factors. These twins can be used as a rich training set for the AI model, which can expand the cohort from dozens to hundreds of thousands. Next the virtual cohort can be filtered to identify patients likely to respond to a treatment, increasing the chances of a successful trial for the target population.

The trial design can also include a sampling of patient types less likely to respond or with elevated risk factors, allowing regulators and clinicians to understand the risks to the broader population without jeopardizing overall trial success. This enhances precision and efficiency in clinical research, providing population-level insights previously available only after years of real-world evidence.

Of course, though today's heart digital twins are powerful, they're not perfect replicas. Their accuracy is bounded by three main factors: what we can measure (for example, image resolution or the uncertainty of how tissue behaves in real life), what we must assume about the physiology, and what we can validate against real outcomes. Many inputs, like scarring, microvascular function, or drug effects are difficult to capture clinically, so models often rely on population data or indirect estimation. That means predictions can be highly reliable for certain questions but remain less certain for others. Additionally, today's digital twins lack validation for predicting long-term outcomes, because the technology has been in use for only a few years.

Virtual twins could demystify the body for patients, fostering trust and encouraging proactive health decisions.

Over time, each of these limitations will steadily shrink. Richer, more standardized data will tighten personalization of the models. AI tools will help automate labor-intensive steps. And the collection of longitudinal data will improve the model's ability to predict how the body will evolve over time.

Throughout modern medicine, new technologies have sharpened our ability to diagnose, providing ever-clearer images, lab data, and analytics that tell physicians what is presently happening inside a patient's body. Virtual twins shift that paradigm, giving clinicians a predictive tool.

Early demonstrations are already appearing in many areas of medicine, including cardiology, orthopedics, and oncology. Soon, doctors will also be able to collaborate across specialties, using a patient-specific virtual twin as the common ground for discussing potential interactions or side effects they couldn't predict independently.

Although these applications will take some time to become the standard in clinical care, more changes are on the horizon. Real-time data from wearables, for example, could continuously update a patient's personalized virtual twin. This approach could empower patients to understand and engage more deeply in their care, as they could see the direct effects of medical and lifestyle changes. In parallel, their doctors could get comprehensive data feeds, using virtual twins to monitor progress.

Imagine a digital companion that shows how your particular heart will react to different amounts of salt intake, stress, or sleep deprivation. Or a visual explanation of how your upcoming surgery will affect your circulation or breathing. Virtual twins could demystify the body for patients, fostering trust and encouraging proactive health decisions.

With the Living Heart Project, we're bringing physics back to physicians. Modern doctors won't need to be physicists, any more than they need to be chemists to use pharmacology. However, to benefit from the new technology, they will need to adapt their approach to care.

This means no longer seeing the body as a collection of discrete organs and considering only symptoms, but instead viewing it as a dynamic system that can be understood and, in most cases, guided toward health. It means no longer guessing what might work but knowing—because the simulation has already shown the result. By better integrating engineering principles into medicine, we can redefine it as a field of precision, rooted in the unchanging laws of nature. The modern physician will be a true physicist of the body and an engineer of health. ■

HOW ARE VIRTUAL TWINS BEING USED IN MEDICINE?

- Virtual twins have guided cardiovascular surgeries, providing predictions and exposing hidden details that even expert clinicians might miss, such as subtle tissue responses and flow dynamics.

- Oncologists are modeling tumor growth and the body's response to different therapies, reducing the uncertainty in choosing the best treatment path for both medical and quality-of-life metrics.

- Orthopedic specialists are personalizing implants to deliver custom-made solutions, considering not only the local environment but also the overall body kinematics that will govern long-term outcomes.

“The Last Human Will See What We Saw”

At 0.016 seconds after the atomic detonation, the fireball was already hundreds of meters wide. The tiny squares to the left in this image are billboards 200 meters from the center of the explosion.

LOS ALAMOS NATIONAL LABORATORY

The test of the world's first atomic bomb left its observers stupefied

By **Emily Seyl**

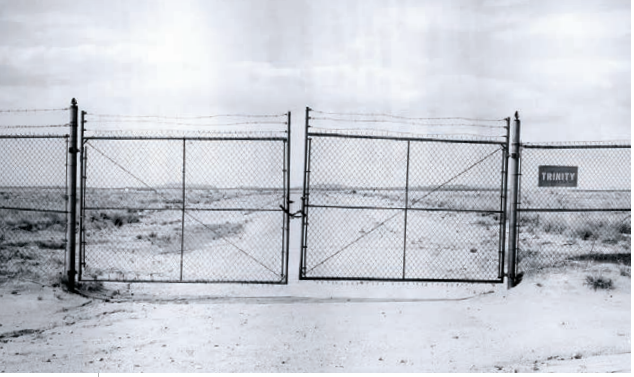

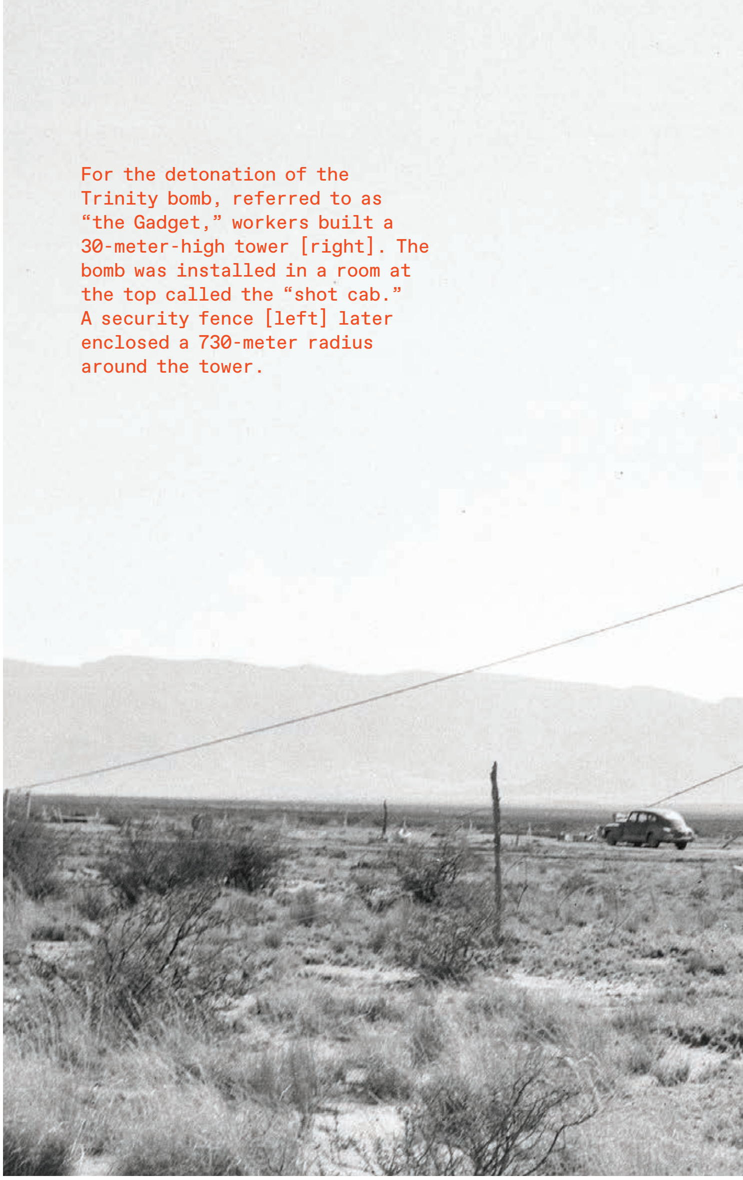

For the detonation of the Trinity bomb, referred to as "the Gadget," workers built a 30-meter-high tower [right]. The bomb was installed in a room at the top called the "shot cab." A security fence [left] later enclosed a 730-meter radius around the tower.

CLOCKWISE FROM BOTTOM LEFT: DAVID WOODFIN/LOS ALAMOS NATIONAL LABORATORY; WHITE SANDS MISSILE RANGE MUSEUM; LOS ALAMOS NATIONAL LABORATORY (4)

Editor's note: If you'd like to pinpoint the instant when the world entered the nuclear age, 5:29:45 a.m. Mountain War Time on 16 July 1945, is an excellent choice. That was the moment when human beings first unleashed the power of the nucleus in an immense, blinding ball of fire above a gloomy stretch of desert in the Jornada del Muerto basin in New Mexico. Emily Seyl's Trinity: An Illustrated History of the World's First Atomic Test *(The University of Chicago Press, 2026) offers hundreds of startlingly vivid photographs of the Manhattan Project that emerged from a 20-year restoration effort. This excerpt and the accompanying photos record the massive effort to capture the awesome detonation of "the Gadget."*

* * *

In the North 10,000 photography bunker, Berlyn Brixner was listening to the countdown on a loudspeaker, his head inside a turret loaded with cameras and film. He was one of the only people instructed to look toward the blast—through his welder's glasses—ready to follow the path of the fireball as it launched into the sky. The two Mitchell

The main photography stations for Trinity were located about 9 kilometers from the shot tower, which was ground zero. The photography group converted machine-gun turrets into sturdy metal tripods [right] that held camera gear and were mounted on the roof of the stations. Some 425 people witnessed the Trinity test. Many of them were issued welder's goggles [above] for safely viewing the blast, because the initial burst of ultraviolet light was predicted to be many times as intense as that of the sun.

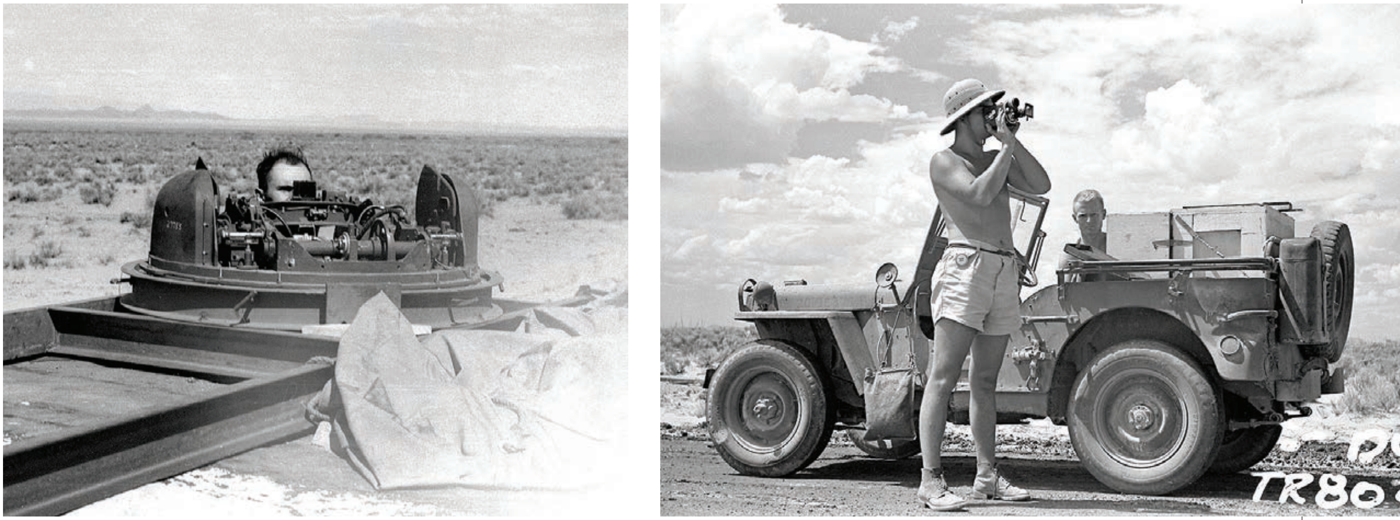

The camera turrets [above] installed on top of several bunkers were spaced at regular intervals around the detonation site. The cameras were activated by a timing signal, while the turrets' positions were controlled by technicians in the bunkers. Ernest D. Wallis [standing, right] was one of dozens of people who photographed Project Trinity. Besides his photographic duties, Wallis was responsible for catching rattlesnakes, which he stashed in the box in the back of his jeep.

The shot tower, which was ground zero for the Trinity test, was in the Jornada del Muerto ("dead man's journey") desert, 340 kilometers south of Los Alamos.

movie cameras at his station would deliver the best footage to come out of the Trinity test, used by Los Alamos scientists to make some of the first measurements of the effects of a nuclear explosion.

When the detonators fired, the cameras captured what Brixner could not have seen—the very first light of a violent, silent sea of energy unfurling into the basin. As 32 blocks of high explosives erupted all together, their incredible force surged inward toward the sleeping plutonium core, compressing the dense sphere of metal instantaneously from all sides and bringing its atoms impossibly close together. A carefully timed burst of neutrons sowed momentary, uncontrolled chaos, and then, as quickly as it began, the fission chain reaction ended. Footage from a high-speed Fastax camera in Brixner's bunker, shot through a thick glass porthole, shows a translucent orb bursting through the darkness less than a hundredth of a second after detonation, as a rush of heat, light, and matter blew apart the Gadget.

When the brightness faded enough for witnesses to make out ground zero, they saw a wall of dust rise up around a brilliant, shape-shifting, multicolored ball of flames—forming a fiery cloud that shot into the sky atop a twisting stream of debris. The camera footage tells a story no less dramatic but hundreds of times more intricate, preserving the moment for scientists to return to again and again to measure and describe the behavior of the fireball and other visible effects with exacting detail. On balance, the photography effort was a huge success, despite only 11 of the 52 cameras producing satisfactory images. By arranging those cameras at intentionally staggered distances, complementary angles, and with a

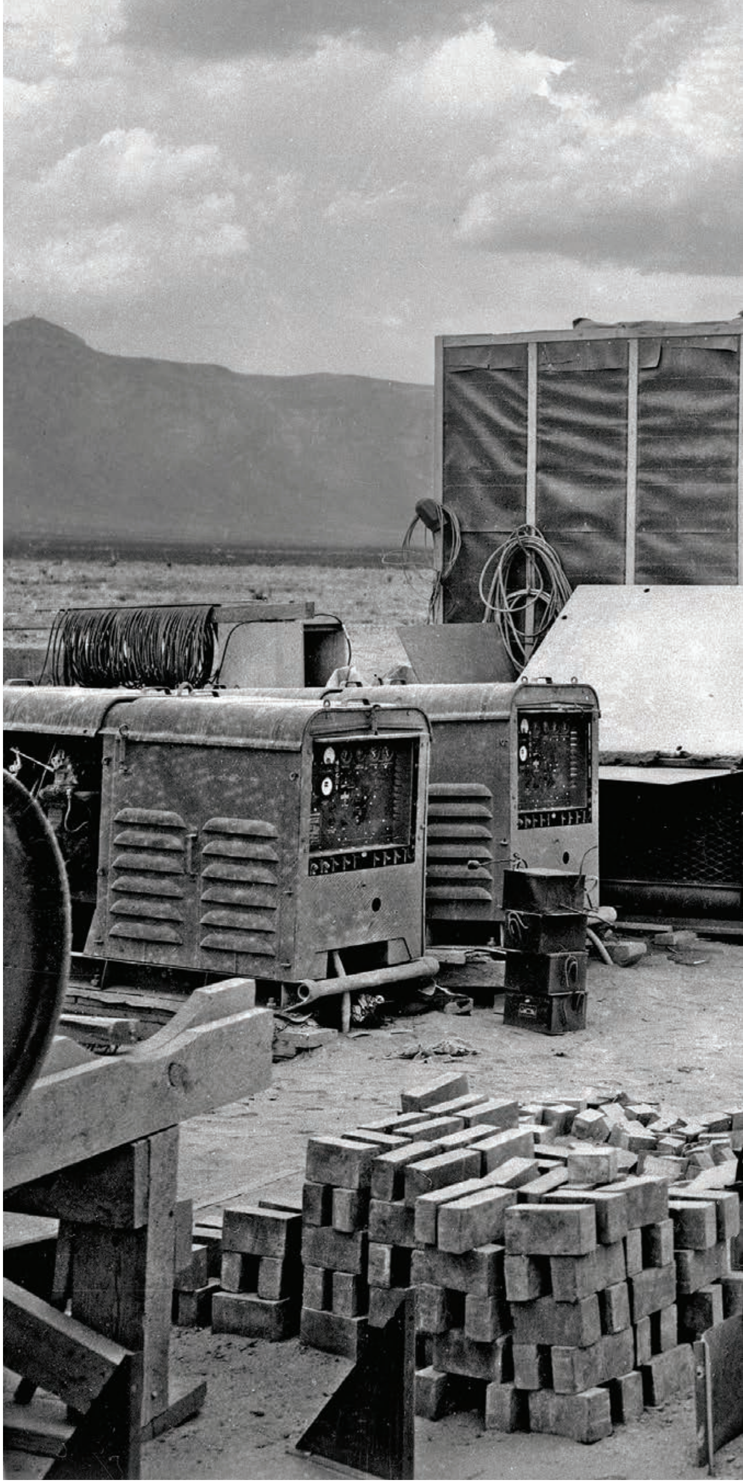

LOS ALAMOS NATIONAL LABORATORY (3)

The North 1,000 instrument bunker [top] was one of several located closest to the shot tower, visible on the right-hand edge of the photo. Trinity technicians used barrage balloons [left] to carry cameras and radiation-measuring instruments into the sky during the blast. One of the instruments, designed by physicist Emilio Segrè, measured gamma radiation in the first fraction of a second after the detonation.

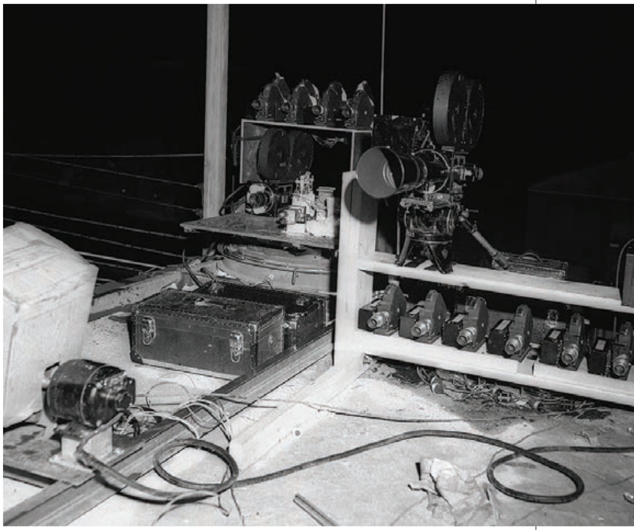

A Mitchell 35-mm movie camera, a military version of the standard camera used in Hollywood in that era, was mounted on the roof of the North 10,000 bunker [above]. On the shelf below the Mitchell camera were 16-mm Ciné-Kodak Model Es, loaded with Kodachrome color film. None of the Kodak cameras, it would turn out, captured good imagery.

Map by Chris Philpot

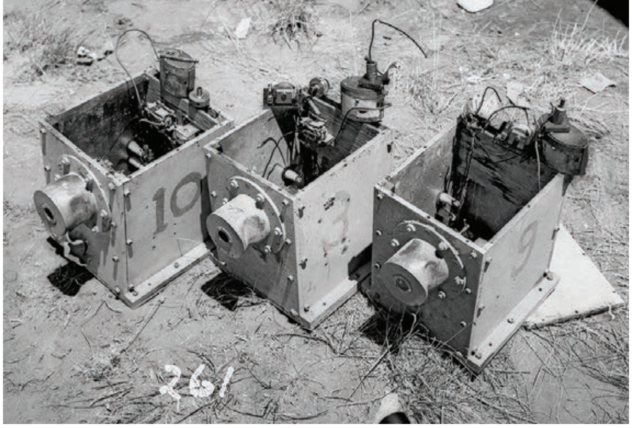

To measure the powerful airborne shock wave that would emanate from the blast, officials used impulse gauges [left], which each contained a water-filled tube with a piston inside. The pressure of the shock wave forced the water out of the tube in the suspended assembly [below, left] and jolted the piston, which moved a stylus that scratched a glass disk. To measure the ground shock waves, buried amplifiers [below] boosted the signals from geophones set up at intervals underground. The signals terminated in equipment racks in the 10,000-yard (9,100-meter) instrument bunkers [right].

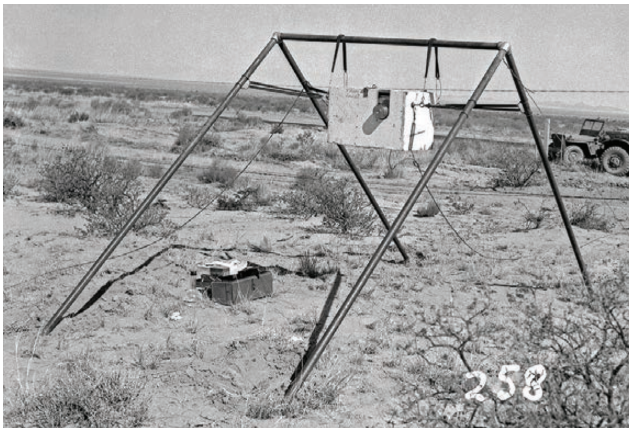

At noon on 15 July 1945, military police closed the road that connected the base camp to the Trinity test site, allowing only authorized personnel through the checkpoint.

On 12 July 1945, Herbert Lehr [above], a U.S. Army sergeant and electrical engineer assigned to Los Alamos, delivered the plutonium core to the McDonald ranch house, where the bomb was assembled. The adobe house, about 3 kilometers from the blast site, had been commandeered by the U.S. military in 1942. Norris Bradbury [right], the physicist responsible for the final assembly of the Gadget, stands next to the partially assembled bomb at the top of the shot tower. The cables on the outside of the bomb would transmit the signals to trigger the synchronized detonations of conventional explosives, which would then create the inward-directed shock wave that would compress the bomb's plutonium core. Bradbury would go on to succeed Robert Oppenheimer as director of Los Alamos on 17 October 1945.

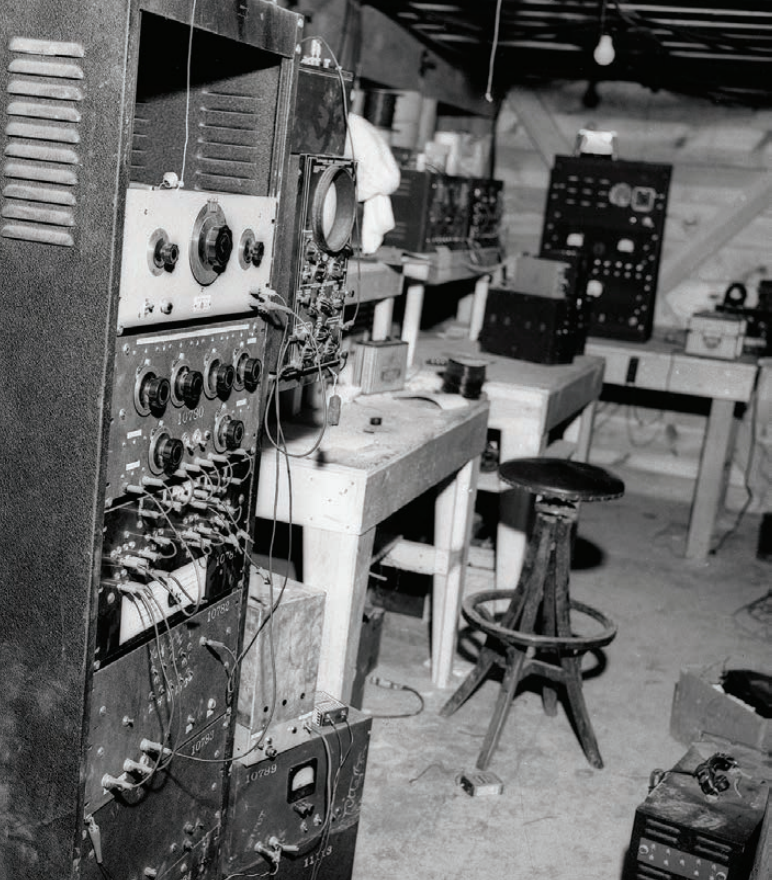

broad spectrum of frame rates and focal lengths, the Spectrographic and Photographic Measurements Group was able to piece together a remarkably complete picture of their subject.

According to the group's leader, Julian Mack, the more than 100,000 frames that were captured still "give no idea of the brightness, or of time and space scales." Mack attributed fortune, as much as foresight, to the photographic record that was made, especially during the earliest phase of the blast. Indeed, the explosion was several times more powerful than predicted, and the intensity of its effects overwhelmed many of the cameras and diagnostic instruments. The human observers were similarly overcome. "The shot was truly awe-inspiring," said Norris Bradbury, the physicist who would succeed Robert Oppenheimer as director of Los Alamos. "Most experiences in life can be comprehended by prior experiences, but the atom bomb did not fit into any preconception possessed by anybody. The most startling feature was the intense light."

It is a common sentiment that words and even pictures pale in comparison to the experience of the explosion. Even so, soldiers, scientists, and many other witnesses have added their firsthand

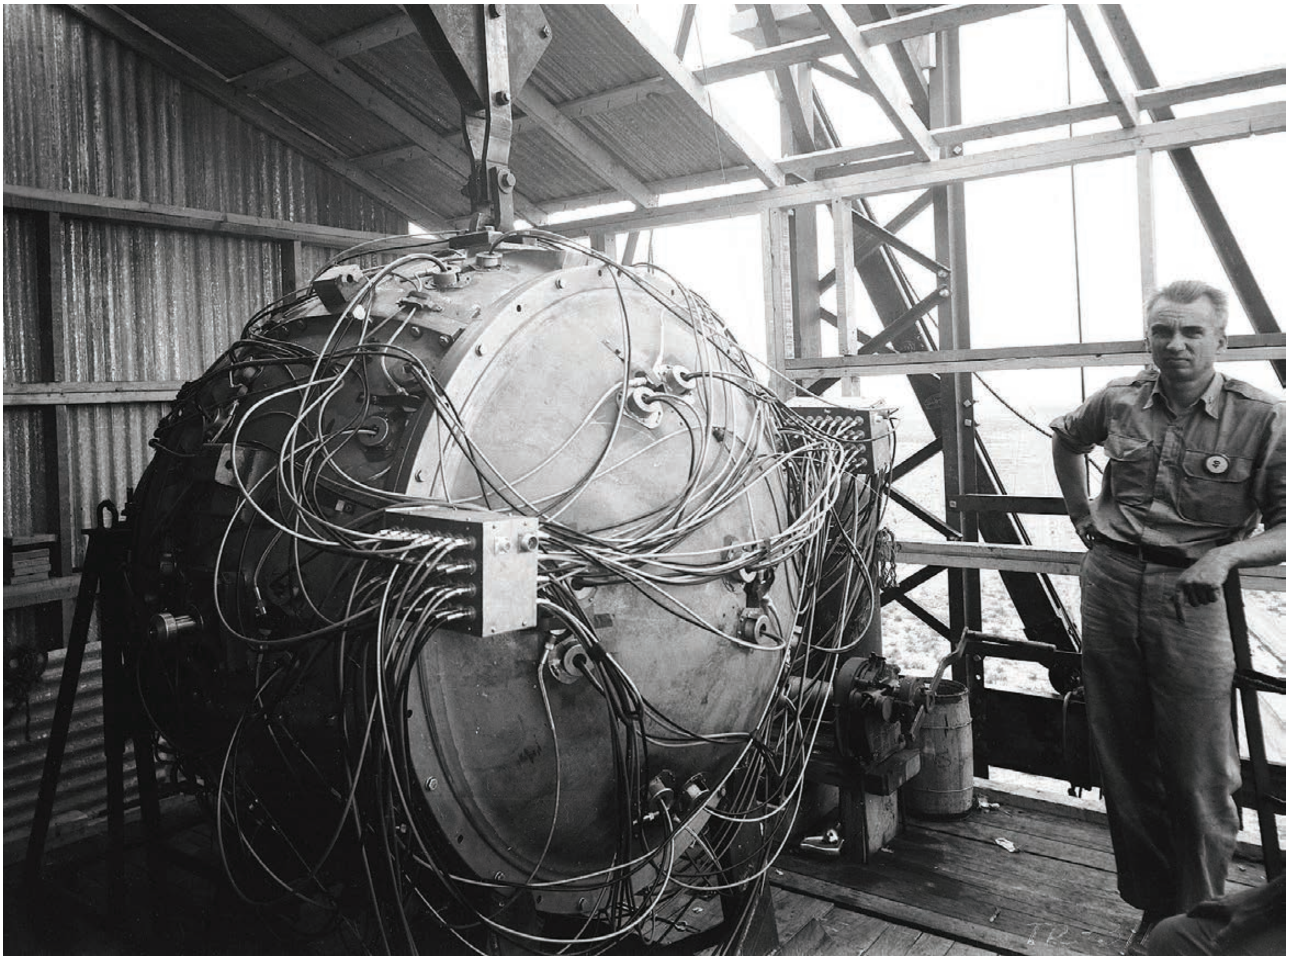

LOS ALAMOS NATIONAL LABORATORY (7)

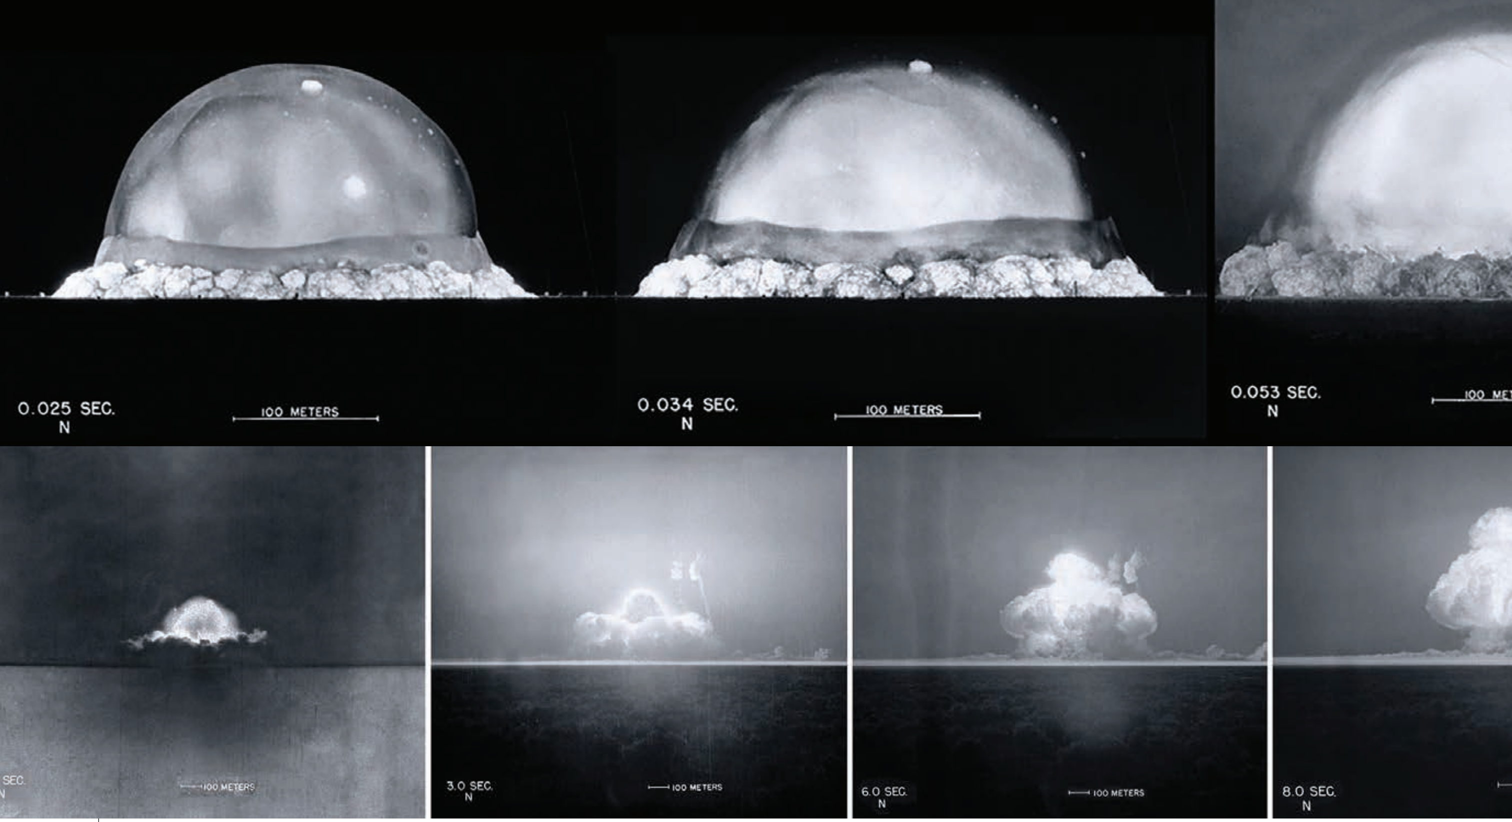

The blast, captured with an assortment of high-speed and motion-picture cameras, shows the fireball expanding between 25 milliseconds [upper left] and 60 seconds [lower right], by which time the mushroom cloud is over 3 kilometers high.

accounts—often absorbing and poetic—to complement the trove of hard data collected during the test shot. They describe an intense and blinding brightness that filled the basin with daytime; an ominous, darkening cloud rearing its head in eerie silence; the wait for the invisible wave rushing out from the heart of the Gadget; and the mighty roar that arrived at last, in a thunder, and seemed never to leave. Physicist Isidor Isaac Rabi, watching from 20 miles away, remembered, “It blasted; it pounced; it bored its way right through you.”

James Chadwick, head of the British contingent of scientists who joined the Manhattan Project, later said, “Although I had lived through this moment in my imagination many times during the past few years and everything happened almost as I had pictured it, the reality was shattering.”

And physicist George Kistiakowsky found himself certain that “at the end of the world—in the last millisecond of the Earth’s existence—the last human will see what we saw.” ■

- - - - -

Reprinted with permission from *Trinity: An Illustrated History of the World’s First Atomic Test* by Emily Seyl with contributions by Alan B. Carr, published by The University of Chicago Press.

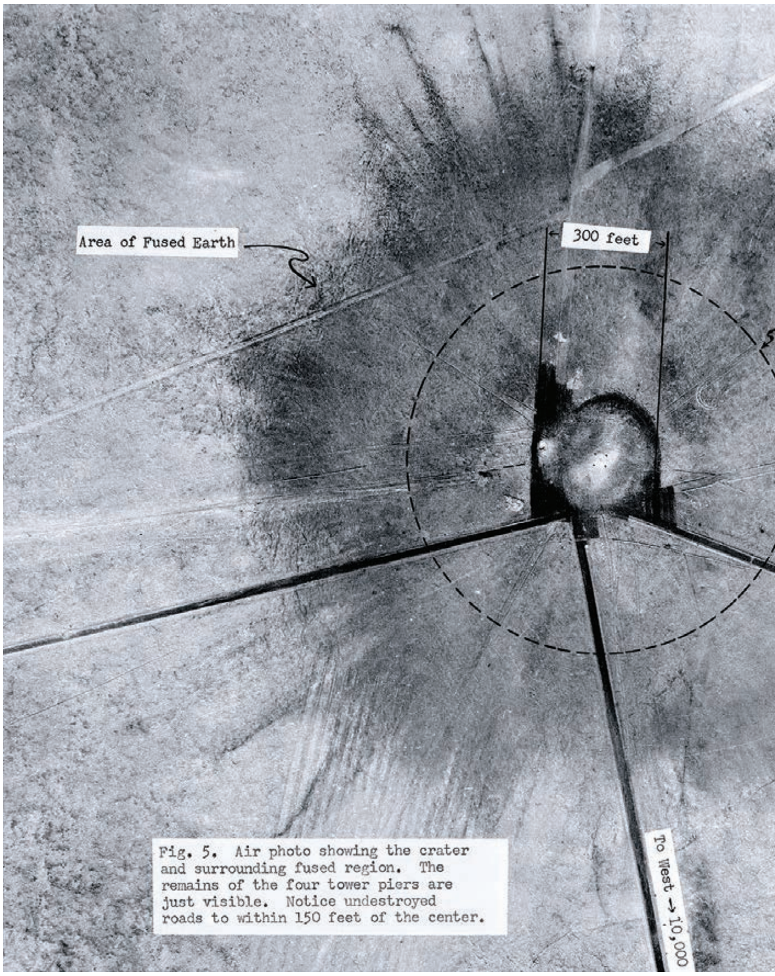

Fig. 5. Air photo showing the crater and surrounding fused region. The remains of the four tower piers are just visible. Notice undestroyed roads to within 150 feet of the center.

CLOCKWISE FROM BOTTOM RIGHT: ETHAN FROGGET/LOS ALAMOS NATIONAL LABORATORY; LOS ALAMOS NATIONAL LABORATORY (14)

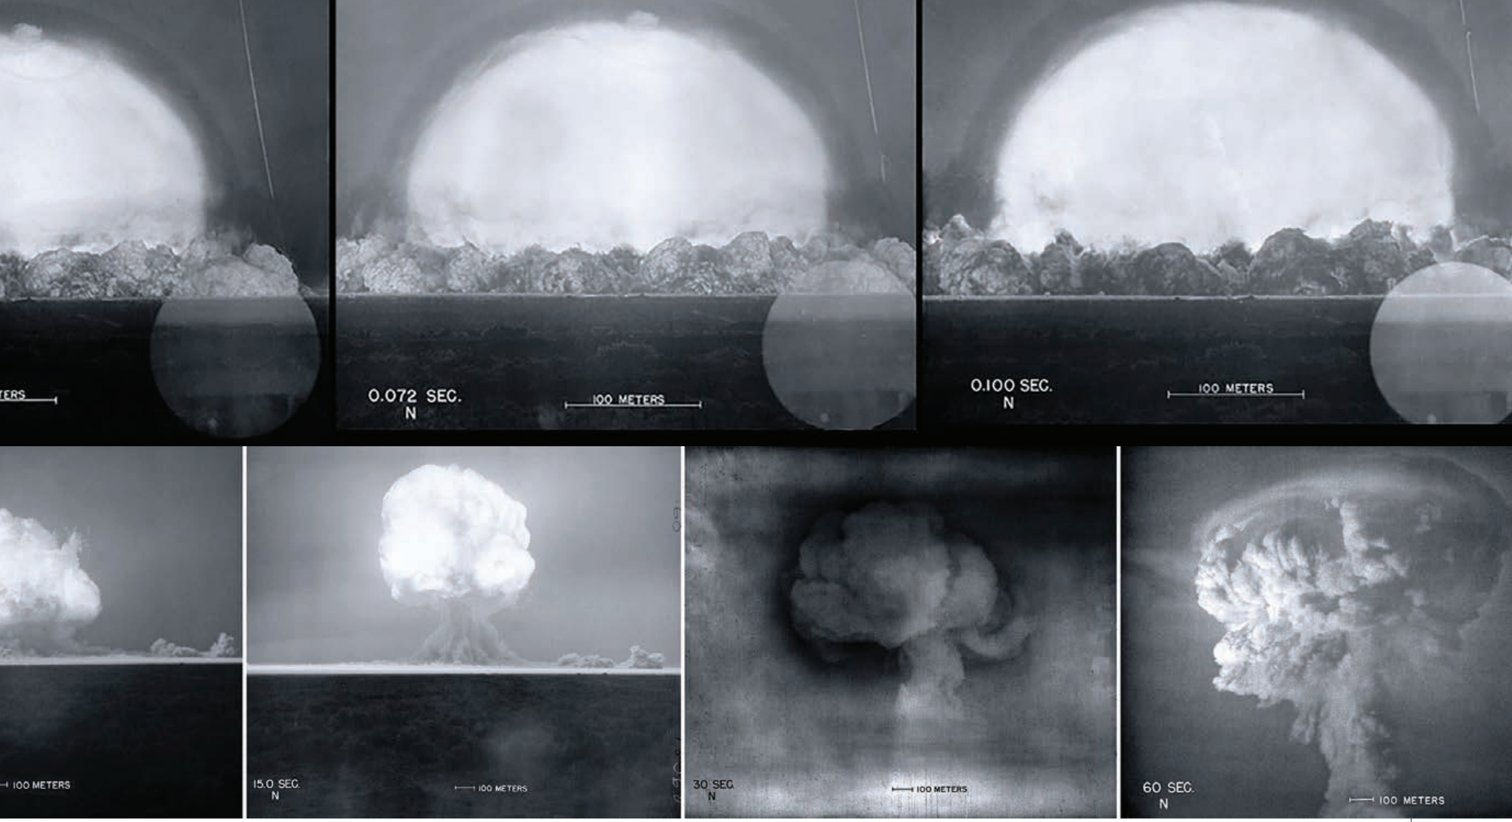

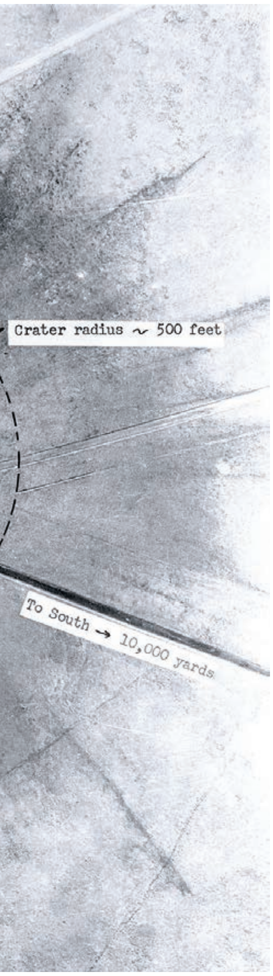

An aerial image [far left] of ground zero shows the postblast distribution of trinitite, formed when sand and bomb and other parts melted during the blast. The trinitite was darker than the surrounding unmelted sand. On 9 September 1945, one week after the end of World War II, Leslie Groves [in khakis and garrison cap], the director of the Manhattan Project, and Robert Oppenheimer, the director of the Los Alamos Laboratory, hosted members of the press at the Trinity site [left]. Since 1965, a lava-rock obelisk [above] has marked the location of the shot tower from the Trinity test.

Sardinia's Energy Future Hinges on Its Past

CONTINUED FROM P. 27

Energy Vault will remove old mining equipment from the Carbosulcis coal mine near Gonnesa to make way for an underground data center [bottom]. It will be powered by a pumped-hydro energy storage system that flows through the mine's vertical geometry and stores water in above-ground tanks [top].

SOMETHING MORE ACCESSIBLE that the public can get behind is building renewables on Sardinia's abandoned industrial sites. "To be honest, not everything is so beautiful here. We have a lot of industrial areas where you can place PV panels. We have a lot of rooftops," electrical engineer Pilo says. "We have unused coal mines." I visit one such project that's proceeding with local support—or at least without much opposition. It's a coal mine near Gonnesa that shut down in 2018 and is now being turned into a data center and a pumped-hydro energy storage system.

Energy storage developer Energy Vault is building it, and despite being based in Lugano, Switzerland—that is, not Sardinia—the company seems to have avoided protest. It helps that the mine is owned by Carbosulcis, a Sardinian regional-government-owned company, which is calling the shots on the project.

Plus, doing nothing with the mine costs money. The mine closed eight years ago because it wasn't profitable, but Carbosulcis must continue maintaining it because of its high methane emissions, which require monitoring and ventilation to prevent explosions and leaks. Carbosulcis managers figured that if they're going to continue putting money and personnel into the mine, they might as well do something useful with it, says Luca Manzella, vice president for Europe, Middle East, and Africa at Energy Vault, as he and I tour the mine.

An innovative project in Sardinia's interior—Energy Dome's grid-scale carbon dioxide battery—seems to be avoiding protest as well. Built in a gated industrial complex near Ottana, this energy-storage facility looks like a giant bubble—the kind that fits over a stadium or tennis complex. It's filled with carbon dioxide that is compressed to store 200 MWh of electricity for the grid. Although the bubble is visible from several of the surrounding hillside villages, and although the developer is headquartered on the mainland, there's little sign of public pushback.

Another path forward is through "energy communities." In this grassroots approach, consumers work together to build their own solar plant or other power generation. Dozens of these communities are already active on the island, according to the Sardinian Electricity Association, a group that provides guidance to consumers.

But by far the greatest need is for energy developers and authorities to understand the people and the history of the land on which they want to build. "When Europe or the national government make a law, they have to also consider the background of Sardinian people and why they are so afraid," says Simone Micheletti, CEO at Futura Group, a renewable-energy developer based in Serramanna, Sardinia. "You cannot apply the same law to Sweden and Sicily. Sometimes you need to understand [the situation] locally," he says.

Decision-makers everywhere would be wise to listen. Otherwise, they may suffer the same fate as their counterparts in Sardinia: despised by locals, delayed by politics, and surprised at how badly it all went. ■

Special thanks to Luigi Avantaggiato for interpreting and additional reporting.

OLD MINE WILL HOUSE DATA CENTER

Despite Sardinia's widespread opposition to wind and solar power, the island hosts some clever renewable energy projects. One of them is an old coal mine, near Gonnesa, being turned into a pumped-hydro energy-storage system that will power an on-site data center. To reimagine the space, Carbosulcis, a mining company owned by Sardinia, has enlisted Lugano, Switzerland–based Energy Vault.

The company plans to use gravity and the mine's vertical geometry to move water through a giant, soft pipe. This will turn a turbine, generating electricity that will be stored in an on-site battery. Power from the grid and a local photovoltaic plant will be used to pump the water back up to the surface where it will be stored in vessels, each about the size of a hot-air balloon. From there, the water is ready to cycle through the system again.

The scheme will power a 60-megawatt data center housed both below and above ground. Twenty percent of the energy will come from water, 60 percent from the battery, and the rest from the grid.

50TH ANNIVERSARY — 1976 TO 2026

The Institute

NEWS OF THE IEEE

VOLUME 50 / ISSUE 2

Career Advice for Engineers, From Engineers

The Institute Celebrates 50 Years **P. 53**

7 Ways You Can Flourish in the Age of AI **P. 56**

President's Column

Designing a Safer Digital World for Kids

CHILDREN BORN AFTER 2013 are the first generation to grow up fully immersed in digital systems, which weren't designed with them in mind. One-third of the world's Internet users are younger than 18, according to UNICEF, yet these systems shaping their daily lives were built for adults. They were optimized for engagement and designed long before people understood how profoundly digital environments influence children.

For engineers and technical professionals, online safety is not an abstract policy debate. It is a design challenge that demands rigor, systems thinking, and ethical foresight.

Governments around the world are also beginning to recognize the problem. Policymakers from across Australia, Brazil, the European Union, Indonesia, and the United States are responding to risks engineers have long understood: Addictive features, inappropriate content, opaque data practices, and algorithmic systems shape user behavior in ways that their creators did not fully predict. For years, technology moved faster than governance. Now governance is trying to catch up.

Global Shift Toward Design Reform

The European Union and the United Kingdom have been among the first to act, embedding age-appropriate digital design into their broader children's rights agenda. Drawing on IEEE expertise and global best practices, Indonesia is the first country in Asia, and Brazil is the first country in Latin America, to adopt age-appropriate design regulation. Australia is aiming to limit access to harmful content and addictive design features through age restrictions on certain platforms. And in the United States, in addition to federal efforts, states including California, New York, and Utah are enacting approaches including age-appropriate design principles.

Across these efforts, a shared realization is emerging. Protecting children online is not simply about filtering content or adding parental controls. It requires rethinking the architecture of digital systems regarding how data is collected, how algorithms make decisions, how interfaces influence attention, and how AI interacts with the developing minds of young users.

Engineers and technical professionals understand that design choices are never neutral. They encode values, incentives, and assumptions. When the user is a child, those choices carry greater weight.

This is where IEEE's work becomes more essential.

Protecting Children Online

For more than a decade, IEEE has been building technical and ethical foundations for safer digital experiences. The first IEEE standard on age-appropriate design in 2021 marked a turning point. It offers a structured, principled approach to designing with children's rights in mind. *The Institute*'s 2022 article "Use a New IEEE Standard to Design a Safer Digital World for Kids" highlights how the standard helps translate those principles into engineering practice.

Today the IEEE Standards Association's (SA) Trustworthy Digital Experiences portfolio provides a practical, technically grounded framework for governments and industry. Spanning ethical design, data

COVER: STUART BRADFORD; THIS PAGE: DEANNA DECKER

governance, algorithmic transparency, and child-focused digital well-being, it has already initiated discussions with government stakeholders around the world. This work helps bridge the gap between engineering realities and policy ambitions.

No single country can solve these challenges alone. Many policymakers lack access to the combined expertise in technology, governance, and children's rights needed to act quickly and effectively. This collaborative effort helps close that gap.

The stakes are high. Without coordinated action, public policy will continue to lag behind technology, leaving children exposed to risks that could have been mitigated through thoughtful design. But with the right frameworks, governments can ensure digital systems respect children's rights, support healthy development, and promote well-being.

IEEE's emerging standards and collaborative technology policy work offer a path forward. By grounding national efforts in evidence-based, rights-aligned design principles, IEEE is helping governments move from reactive regulation to proactive, coherent, and globally informed strategies for protecting children online.

Safeguarding childhood in the digital age is both a moral imperative and an engineering challenge. And IEEE is helping to lead the way.

—MARY ELLEN RANDALL
IEEE president and CEO

Please share your thoughts with me: president@ieee.org.

Supporting National Digital Ambitions

In Athens this year I met with senior leaders of Greek government agencies and key national research institutions. Greece is moving quickly on digital transformation and responsible technology governance, and our discussions reinforced IEEE's role as a trusted, neutral collaborator.

We focused on supporting Greece's ambitions in digital modernization and public-sector innovation. We also discussed responsible AI and age-appropriate digital design in Europe and elsewhere. These engagements, grounded in shared values and long-term commitment, strengthened IEEE's presence within the European ecosystem and opened new pathways for collaboration on trustworthy AI and child-focused digital well-being.

From left: Alpesh Shah, managing director of IEEE Standards; Mary Ellen Randall, 2026 IEEE president and CEO; Georgios Nounesis, chair of the National Council for Research, Technology, and Innovation; and Konstantinos Karachalios, former managing director of IEEE Standards.

LEFT: IEEE; RIGHT: NICOLE MILLMAN

Editor's Note

The Institute Celebrates 50 Years

BY KATHY PRETZ

THE INSTITUTE IS celebrating its 50th anniversary this year. Launched in 1976, the publication was designed to keep members informed about IEEE and what its constituents were doing, as well as to report on the organization's initiatives, technical standards, products, and services.

That directive expanded over the years to include our reporting on key historical technical achievements recognized as IEEE Milestones and support for young professionals with career-guidance articles and information about educational resources.

The Institute has gone through many iterations in the past 50 years.

Editor's Note

What began as a monthly four-page insert in the print edition of *IEEE Spectrum* became a separate newspaper published six times a year and mailed along with *Spectrum* in 1977, and then a monthly publication the following year.

Today we publish all of *The Institute*'s articles online, with a curated selection appearing in our 16-page quarterly, printed in the March, June, September, and December *Spectrum* issues.

To provide a quick summary of the latest online news, a bimonthly e-newsletter, *The Institute* Alert, began appearing in members' inboxes in 2003. Members also can stay up to date by following our Facebook, Instagram, and LinkedIn pages.

Although much has changed, an original subsection from 1976—"IEEE People"—has been maintained for the past five decades. We continue to celebrate IEEE members from around the world through our profiles, which are among our most popular articles.

As the longest-serving editor in chief for *The Institute*, I count it a privilege that my staff and I get to chronicle the stories of remarkable IEEE-affiliated individuals. They are usually unseen visionaries and problem-solvers who work tirelessly behind the scenes on technologies that are reshaping the world.

By highlighting their careers and how IEEE has played a role in their professional growth, we hope to inspire the next generation of engineers and technologists to continue a legacy of innovation and service to humanity. ■

Learn In-Demand Skills With IEEE Programs

BY ANGELIQUE PARASHIS

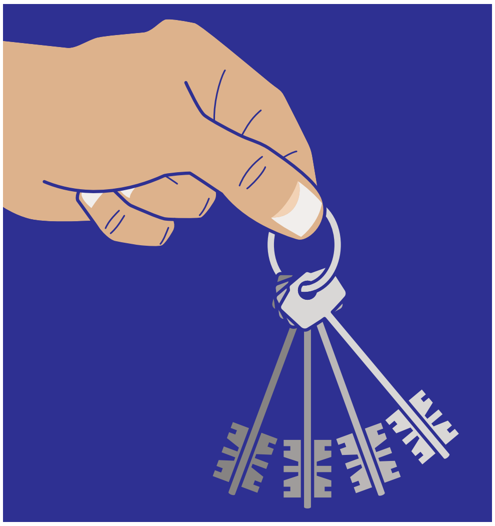

TECHNICAL PROFICIENCY DOES not guarantee organizational leadership acumen. To bridge the gap, the IEEE Professional Development Suite offers training programs designed to build the strategic competencies required to navigate today's complex landscape.

The suite, developed by IEEE Educational Activities, provides deep technical dives into domains such as telecommunications connectivity and microelectronics reliability.

Electrostatic Discharge and 5G

Electrostatic discharge (ESD) is a major reliability challenge. The online Practical ESD Protection Design certificate program covers theory, real-world case studies, and practical mitigation techniques.

As 5G network capabilities expand globally, so does the demand for engineers. The IEEE 5G/6G Essential Protocols and Procedures Training and Innovation Testbed, in partnership with British company Wray Castle, takes a deep dive into the 5G function framework, registration processes, and packet data unit session establishment.

The program is designed for system engineers, integrators, and technical professionals responsible for 5G signaling.

Leadership Training

Technical knowledge alone is not enough to climb the corporate ladder. Engineering leaders must have strategic vision and people-centric leadership skills.

The IEEE Leading Technical Teams training program focuses on the challenges of managing engineers in R&D environments and fostering creative problem-solving through an immersive learning experience. It's designed for professionals who have been in a leadership position for at least six months.

Organizations can enroll groups of 10 or more to learn as a cohort.

IEEE, in collaboration with the Rutgers Business School in New Jersey, offers two mini MBA programs to bridge the gap between technical expertise and executive leadership. During the mini MBA for engineers 12-week curriculum, technical professionals master core competencies such as financial analysis, business strategy, and negotiation so they can effectively transition into management roles.

The mini MBA in artificial intelligence embeds AI literacy directly into business strategy. Participants learn to evaluate AI through financial modeling and governance frameworks.

All the programs offer continuing education units and professional development hours. ■

Angelique Parashis is the senior manager of education marketing for IEEE Educational Activities.

Illustration by MCKIBILLO

Improve Your Writing Skills With This IEEE Course

BY ANGELIQUE PARASHIS

ONE OF THE most important skills an engineer can develop is the capacity to write and communicate effectively.

Writing is often labeled a “soft skill”—which can diminish its importance. In reality, communication is a core engineering competency. It lets you document methods, articulate research findings, and persuade decision-makers who determine whether projects move forward.

If your writing is dense, disorganized, or overloaded with technical jargon, the value of the underlying work can become obscured. Clear writing can strengthen the impact of your work. Poor writing can distract from the points you’re trying to make, as readers might not understand what you’re saying.

Technical writing differs from other forms of prose because readers expect information to follow predictable, logical patterns. Unclear writing can leave readers unsure of the author’s intent.

Despite technical communication’s importance, engineering curricula often limit or lack formal instruction in it.

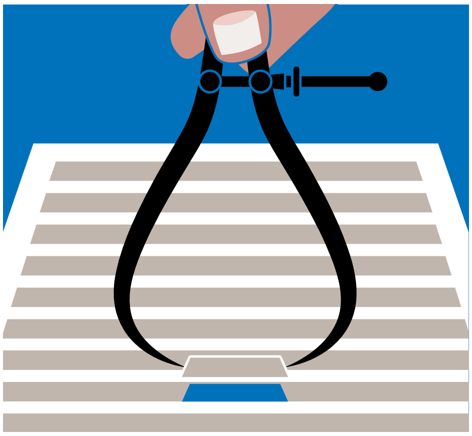

Recognizing that gap, IEEE has expanded its role as a global knowledge leader by offering From Research to Publication: A Step-by-Step Guide to Technical Writing. The course is led by Traci Nathans-Kelly, director of the engineering communications program at Cornell.

The learning opportunity, developed by IEEE Educational Activities and the IEEE Professional Communication Society, goes beyond foundational writing skills. It addresses today’s challenges, such as the ethical use of generative AI in the writing workflow, the complexities of team-based authorship, and publishing strategies.

The program centers on core skill areas that can influence an engineer’s ability to communicate. Participants learn to master the IMRaD structure and learn advanced editing techniques to help strip away jargon so complex ideas are more accessible. In addition, the course covers strategic approaches to publishing work in high-impact journals and improving a writer’s visibility within the technical community.

The course is available on the IEEE Learning Network at iln.ieee.org. Participants earn professional development credit and a shareable digital badge. IEEE members receive a US $100 discount. ■

Angelique Parashis is the senior manager of education marketing for IEEE Educational Activities.

Designing Microcredential Programs With Universities

THE RAPID ASCENT of artificial intelligence and semiconductor manufacturing has created a paradox: Industries are booming yet they face a critical shortage of skilled workers. Demand for data center technicians, fabrication facility workers, and similar positions is growing. Although those technical roles are essential, they don’t always require a four-year degree—which has paved the way for skills-based microcredentials. By partnering with higher education institutions and training providers, industry leaders are helping to design targeted skills programs that quickly turn learners into job-ready technical professionals.

Skills Validation

Through its credentialing program, IEEE serves as a bridge between academia and industry. The program, developed and managed by IEEE Educational Activities, offers standardized credentials in collaboration with training organizations and universities.

IEEE is setting the benchmark for skills-based microcredentials by establishing a framework that includes assessment methods, qualifications for instructors and assessors, and criteria for skill levels.

A collaboration with the University of Southern California, in Los Angeles, for example, developed microcredentials for USC’s

Illustration by MCKIBILLO

semiconductor clean room program. Standardized skills assessments and associated microcredentials were created to help people join the semiconductor industry as clean room technicians or as engineers with cleanroom experience.

IEEE also has partnered with the California NanoSystems Institute at the University of California, Los Angeles, to create skills-based microcredentials for its clean room protocol and safety program.

Best Practices

Based on IEEE's work designing microcredentials, three best practices have emerged.

1. Align with industry needs before design. Collaborate with industry prior to starting the design process. Higher education institutions and training providers build relationships with companies to create effective microcredential programs and methods of assessment.

2. Build for flexibility. An adjustable skills-based microcredentials framework allows programs to create or pivot as new breakthroughs occur.

3. Implement a continuous-feedback loop. Emerging technical fields are still being developed or are quickly evolving. This requires continual communications and feedback among higher education, training providers, and industry. Generating consensus on the skills sets learners must gain is critical for microcredential programs. Setting minimum standards still allows providers to update assessments as new tools and safety protocols enter the workplace.

—Kaitlyn Ridel

Kaitlyn Ridel is the IEEE microcredentials marketing specialist.

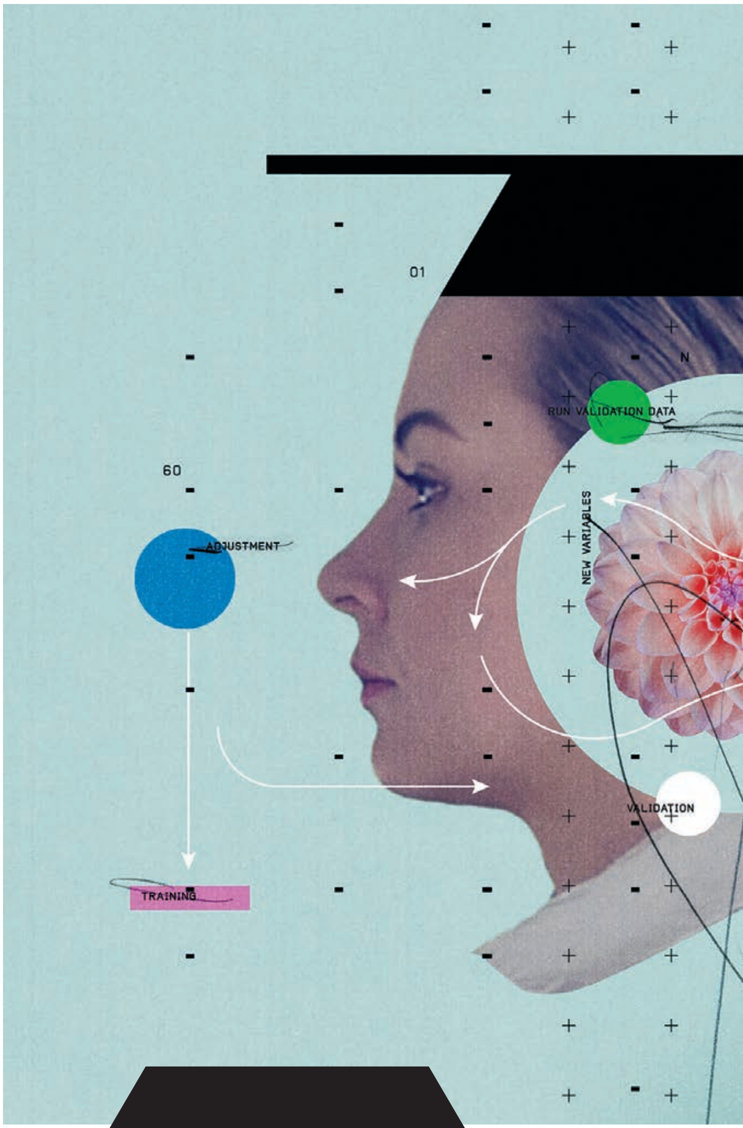

7 WAYS NEW ENGINEERS CAN FLOURISH IN THE MODERN AGE

View AI tools as a complement, not competition

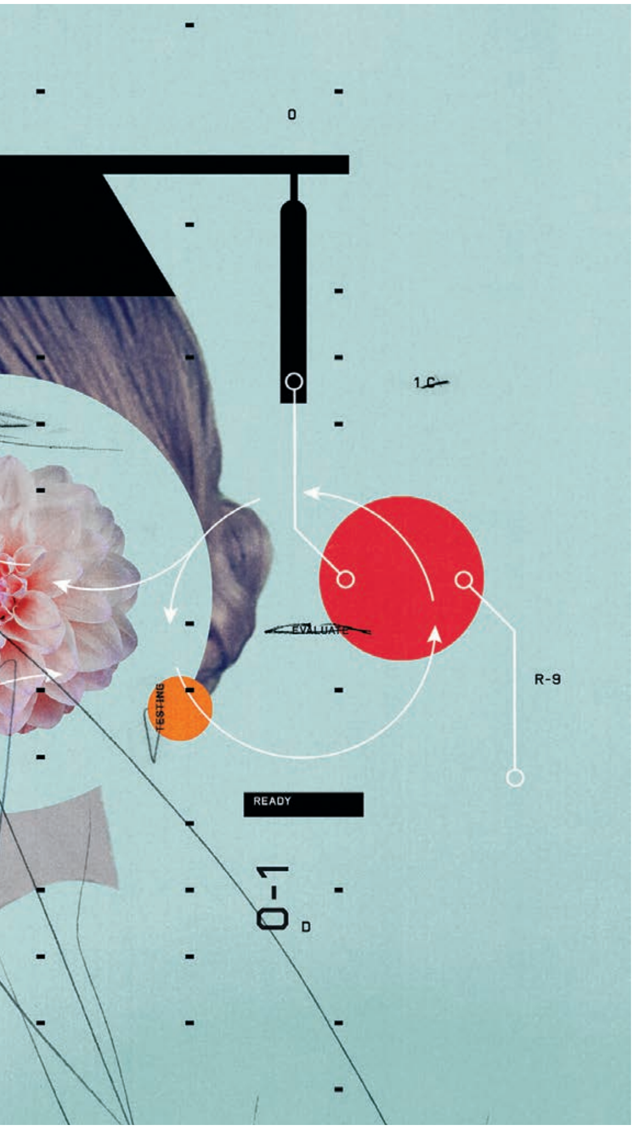

BY LOKESH LAGUDU

NEW GRADUATES' CAREERS are unfolding in an era when AI is not optional. The most successful engineers treat artificial intelligence as leverage, not competition.

Here are seven tips to help keep young professionals in demand.

1. Master the fundamentals first. AI tools can help you code, but you still need strong fundamentals:

- Data structures and algorithms for problem-solving.
- Operating systems, databases, and networking for system-level understanding.
- Core programming languages such as C++, Java, and Python.

AI can autocomplete syntax, but if you don't understand how things work under the hood, you're likely to struggle to debug or optimize.

2. Learn how to work with AI, not against it. The best engineers will not try to out-code AI. Instead, they will:

- Write clear prompts to generate better code snippets.
- Review and debug AI-generated code for accuracy, performance, and security.
- Use AI for productivity boosts while still exercising judgment.

Think of AI as a teammate. The real skill is knowing when to trust it and when not to.

3. Build projects that showcase end-to-end thinking. Employers increasingly look for engineers who can design and build systems, not just solve problems. Create projects that show you can:

- Define requirements clearly.
- Use AI tools responsibly within the workflow.
- Deliver a product that scales and is maintainable.

4. Sharpen your system design skills early. Even junior engineers are now asked questions about basic system design with AI. Expect to explain to prospective employers:

- How you would responsibly integrate AI into a system.
- How to design fallbacks when AI fails.
- How to ensure scalability and reliability.

Continued on page 66

Illustration by Stuart Bradford

HOW TO GAIN YOUR FOOTING IN AI

An IEEE guide offers tips on thriving amid ever-changing tech

BY WILLIE D. JONES

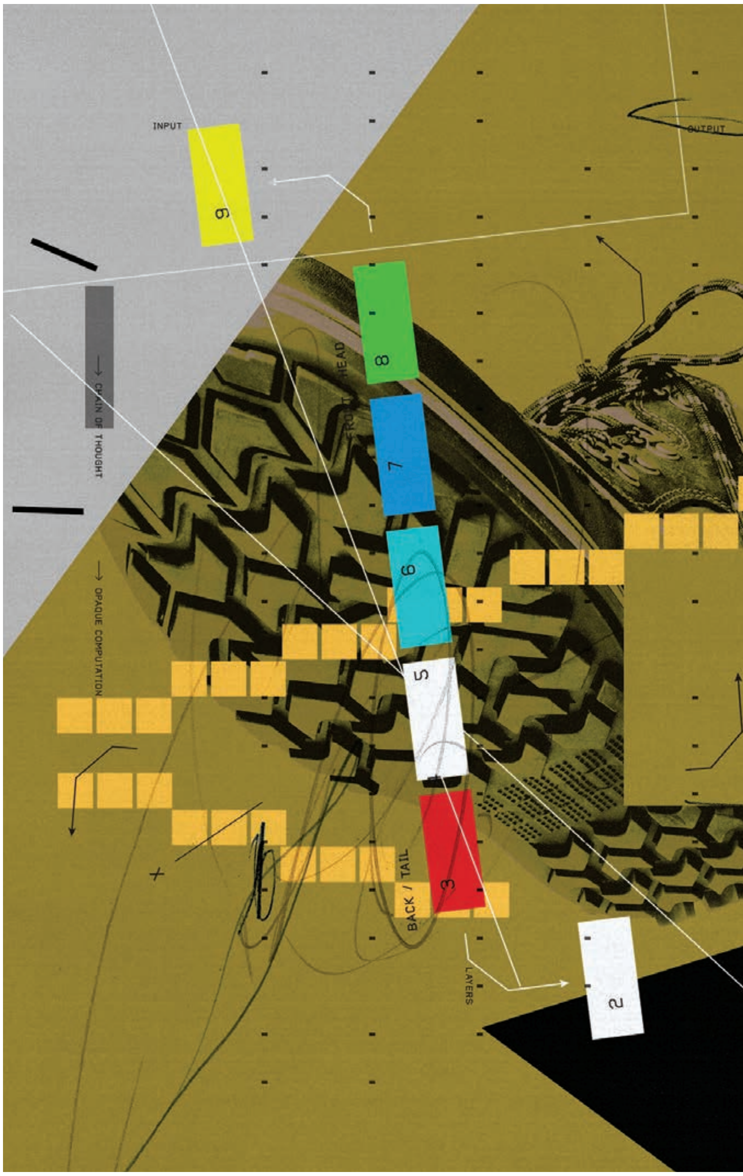

THE NEWLY RELEASED guide "Preparing for a Career as an AI Developer" from the IEEE Computer Society argues that the most durable path into artificial intelligence is not defined by mastering any single tool, programming language, or model. Instead, long-term success depends on combining strong technical foundations with human-centered skills that machines are unlikely to replace.

AI is reshaping the job market faster than academic programs and employers can adapt, according to the guide. Systems now analyze cybercrime, predict equipment failures in factories, and generate text, software code, and images at scale. The rapid advance has contributed to mass layoffs across parts of the technology sector while unsettling recent graduates and early-career professionals entering an uncertain labor market.

Yet demand for AI expertise remains strong. Banking, health care, retail, and pharmaceutical companies are racing to deploy generative AI tools to improve productivity and decision-making, creating a sustained need for workers who understand both the technology and its applications.

The resulting contradiction—job disruption alongside growing opportunity—has left many unsure how best to prepare for careers in a field constantly redefining itself. Addressing that uncertainty is the goal of the guide, written by San Murugesan and Rodica Neamtu. Murugesan is an IEEE life senior member and adjunct professor at Western Sydney University, in Australia.

Illustration by Stuart Bradford

Neamtu is an IEEE member, and a professor of teaching and a data-mining researcher at Worcester Polytechnic Institute, in Massachusetts.

The downloadable 24-page PDF report outlines which skills aspiring AI professionals should prioritize and explains why AI careers increasingly center on applying technology thoughtfully rather than building algorithms in isolation.

Redefining AI Careers

AI's expanding reach distinguishes the current technological moment from earlier waves of automation, the authors write. Pattern recognition, reasoning, optimization, and machine learning now influence nearly every sector of the economy.

The guide cites executive surveys, including research from McKinsey & Co., showing persistent shortages in advanced IT and data analytics talent. Employers also report gaps in critical thinking, creativity, and problem-solving—capabilities that remain difficult to automate.

The authors frame this mismatch as an opportunity for workers who prepare strategically by developing complementary skills rather than competing directly with AI systems.

Neil Thompson, director of FutureTech research at MIT's Computer Science and Artificial Intelligence Laboratory, reinforces that approach in the guide. Workers, he says, should focus on applying AI within adjacent domains where human expertise adds context and judgment.

"When we see rapid technological progress like this, workers should focus on skills and occupations that apply AI to adjacent domains," Thompson says, highlighting scientific research as an area with strong potential.

Fundamentals Still Matter

Adaptability does not replace technical rigor. A viable AI career still requires a solid grounding in computing, data, and machine learning principles.

Core knowledge areas include data structures, large-scale data handling, and tools for data manipulation and analysis. Foundational concepts such as supervised and unsupervised learning, neural networks, and reinforcement learning remain essential.

Because modern AI systems depend on scalable infrastructure, familiarity with cloud platforms—including Amazon Web Services, Google Cloud, and Microsoft Azure—is increasingly important. Mathematics underpins all of it, particularly linear algebra, calculus, and probability theory, which form the basis of most machine-learning algorithms.

Python has emerged as the dominant programming language for experimentation and model development due to its extensive ecosystem of AI libraries.

Hands-on experience with development frameworks also plays a key role. PyTorch, developed by Meta AI, is widely used for prototyping deep-learning systems, Scikit-learn supports classification and regression tasks, and TensorFlow, a software library for machine learning, enables building and deploying machine-learning systems at production scale.

The authors emphasize that such tools matter less as résumé keywords than as ways to understand how models behave under real-world constraints, including imperfect data and operational limits.

Maintain Soft Skills

As AI projects grow more interdisciplinary, human skills are becoming central to technical work. Professionals must explain system behavior, limitations, and risks to nontechnical stakeholders, making communication increasingly important.

Neamtu describes communication and contextual thinking as timeless abilities that gain value as automation expands, particularly when paired with leadership, resilience, and continuous learning. Murugesan similarly stresses that technical depth must be matched by collaboration and adaptability.

Stay Curious

Murugesan urges aspiring professionals to treat continuous learning as a career strategy.

"Always be curious," he says. "Learn from failure. Mistakes and setbacks are part of the journey."

Neamtu adds that because AI will affect nearly every profession, long-term success depends less on chasing specific technologies than on sustaining curiosity and aligning technical skills with meaningful applications.

In a field where tools can become obsolete within a year, the guide's central message is clear: The most future-proof AI career is built not on what you know today but on how well you continue learning as technology changes. ■

WAYS TO BOOST YOUR ODDS OF LANDING YOUR FIRST JOB

Impress employers with real-world projects

BY RAMNEEK KALRA

AS A COLLEGE student, are you concerned that your knowledge alone won't be enough to impress potential employers? Do you feel you lack the necessary hands-on technical skills to secure a job? Maybe you've thought of an engineering solution for a problem in your school or community but are unsure how to take the next step.

I struggled to bridge the gap between classroom theory and real-world application. But when you combine academic knowledge with practical projects that solve a societal problem with technology, you can ace any interview.

You don't have to navigate the journey alone. Here are some lessons I learned as a student.

Speeding Up Processes

I'm a cloud support engineer at a company in Hyderabad, India. I'm also an active IEEE volunteer as one of its young professionals, impact creators, and brand ambassadors.

When I was in my first semester as a computer engineering student at Guru Gobind Singh Indraprastha University, in New Delhi, I became frustrated by the long lines to check books in and out of the library of the affiliated college, the HMR Institute of Technology and Management. Even getting a new library card took a long time. I was determined to solve the problem.

For six months, I singlehandedly developed a software program to scan student ID cards and speed up the processes. Word got out about my programming skills, and I received many requests to help solve other problems.

An intriguing one was from the director of India's largest national broadcasting company, All India Radio. I was asked to streamline its accounting process. At the time, the company used only Microsoft Excel along with a pen-and-paper system. It took me just six months to build a full-stack accounting software program to make the process significantly more efficient.

That opportunity was a big break for me. The technology I created redefined the broadcaster's operations and could be used in its other offices, expanding my reach.

In my first corporate job interview after graduating from university, the interviewer was surprised to learn that I'd published 15 research papers, completed 15 projects, and even had a pending patent application. (The government has since granted the patent.)

Students need to understand the importance of doing something exceptional beyond learning theory and concepts. Having practical skills before leaving school is a great way to set yourself apart from other new engineering graduates.

Ask the Right Questions

Before taking on any new projects, I ask myself five simple questions.

1. Who uses the current model?
2. What are its features?
3. Why is the current model insufficient?
4. When is the right time to deploy the new solution?
5. Where should it be deployed?

Once I get those answers, I begin using design thinking to strategize.

Five Stages

Here are what I consider to be the five stages of the process.

Illustration by Stuart Bradford

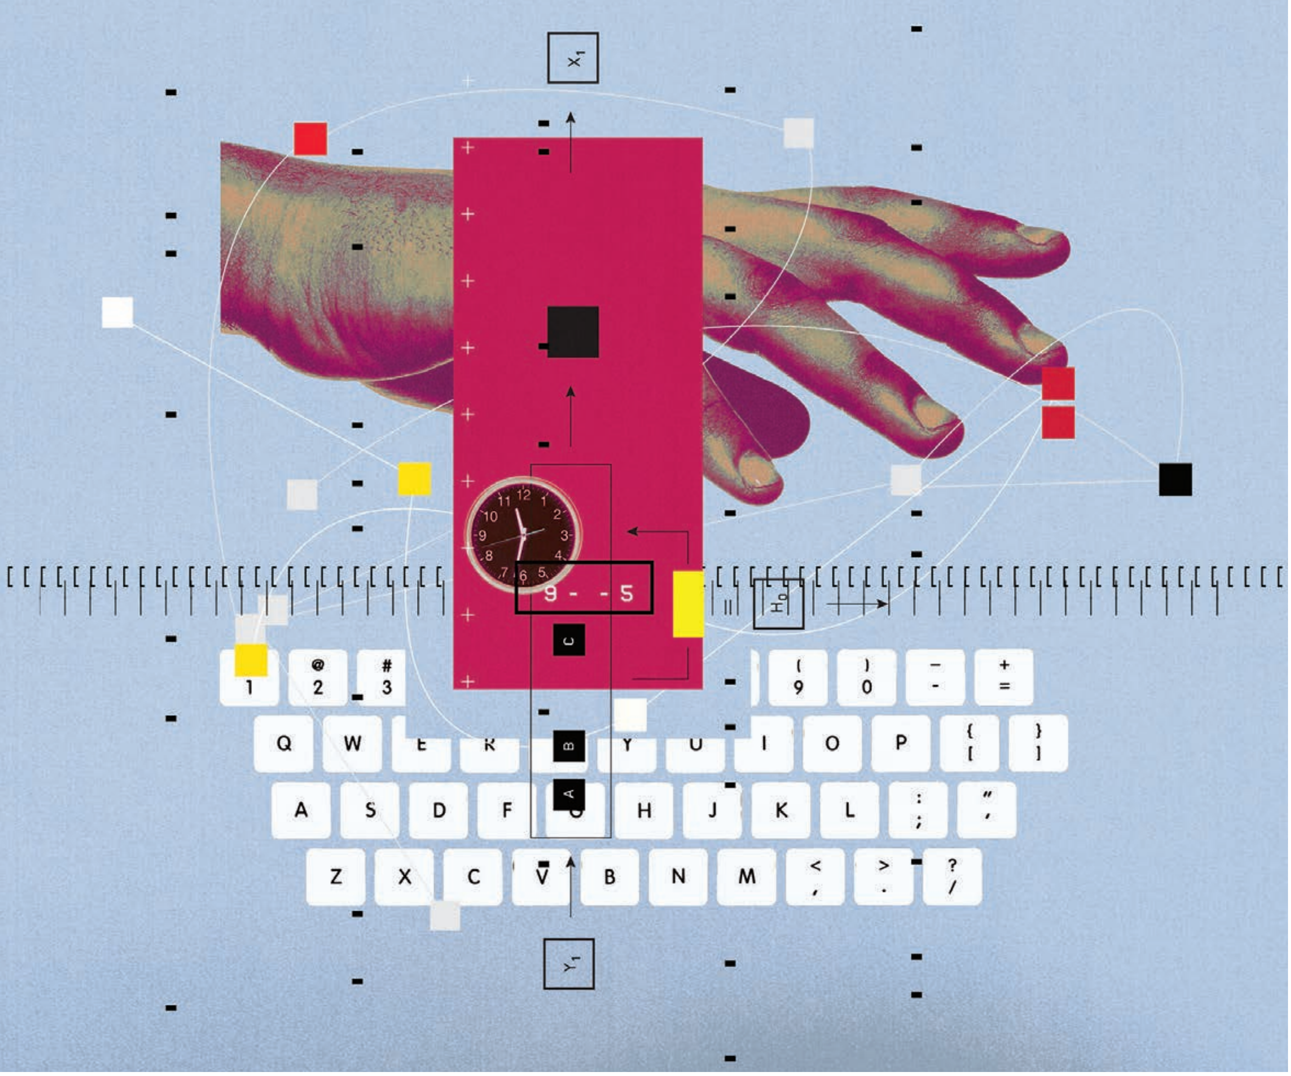

Understand the problem. Once you identify the client's issue, the next step is to listen to them in full without making judgments. You need to really understand the pain points and why the current application isn't working. Listen fully, ask questions, and try to empathize with their issues.

Research and ideation. Do your own research. It's essential to conduct field research to better understand the client's requirements. Next, start brainstorming. Consider how you can improve the current model. Maybe you should conduct research to find other products that might solve the problem. Also consider redesigning the current version.

Technology research and prototyping. By this stage, you've created a short list of ideas to address your client's pain points. Next, research all the technologies you need to use. If you need training, use learning platforms such as Coursera, EdX, the IEEE Learning Network, Udacity, and Udemy.

Once you identify and learn the technology needed, it's time to create the first prototype.

Test and improve. Test the prototype, gather feedback from your client while you take meticulous notes, and then revise it according to the feedback.

That helps you understand what improvements are needed and helps identify gaps in your model. It gets you closer to the client's requirements. Use the information to refine the design and build the product.

It is important to note that this stage might go through multiple iterations. You might have to continue to improve the results until the design works for the client.

Protect your intellectual property. Many students and young professionals skip the important step of safeguarding their idea such as copyrighting it, publishing a paper, or filing a patent. I have seen many students who present their ideas at hackathons and competitions and assume that receiving cash prizes is enough to list on their résumé. They really should protect their ideas. ■

Ramneek Kalra is an IEEE Young Professional, IEEE brand ambassador, and an IEEE impact creator.

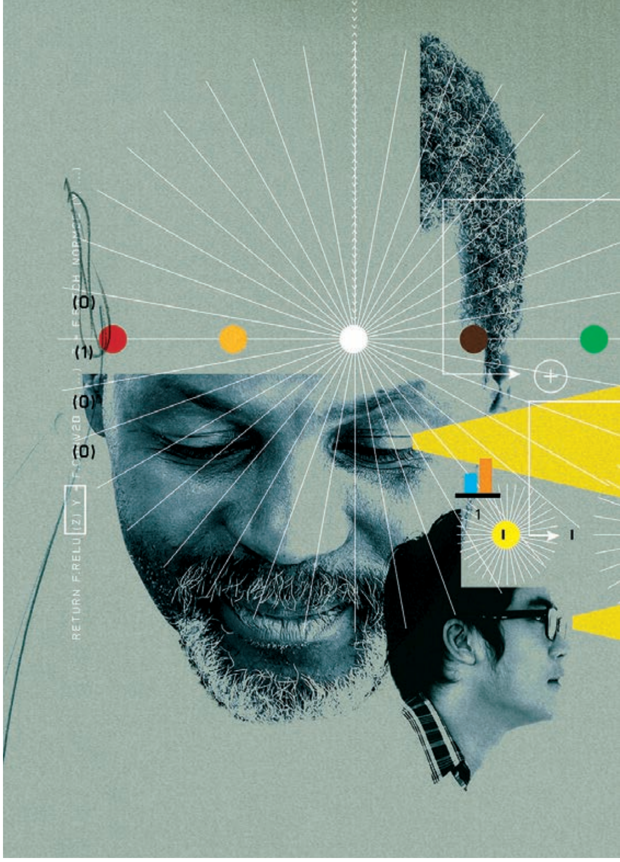

MENTOR YOUR STAFF FOR GREATER INNOVATION

Organizational success builds on individual accomplishments

BY LOKESH LAGUDU

IN THE MODERN era of rapid digital transformation, engineering leaders are expected to be more than project managers and technical experts. They need to be vision-setters, innovation enablers, and mentors shaping the next generation of talent.

Leadership and mentorship, when paired with intention, do more than advance business goals. They create an ecosystem where innovation flourishes and careers accelerate.

Leadership: An Innovation Engine

Innovation rarely happens by accident. It is cultivated in environments where leaders articulate a compelling vision, empower their teams to experiment, and then remove obstacles that stifle creativity.

As a senior engineering manager at Walmart Global Tech in Sunnyvale, Calif., I have led efforts to address the challenge of shrinkage. This loss of inventory, commonly due to shoplifting, theft, and return fraud, results in a difference between the amount of stock a retailer is supposed to have and the amount it actually has. Globally, retailers lose more than US $100 billion annually.

The scale of the problem demands more than incremental improvements. By aligning the challenge with cutting-edge technologies such as computer vision and artificial intelligence, I framed a plan that transformed a business imperative into a technological frontier. We focused on deploying computer vision models at the store front-end, supported by an edge and cloud pipeline that allowed rapid experimentation. The system combined real-time detection of high-risk events with predictive analytics that highlighted emerging patterns of loss, and it integrated directly with store operations so actions could be taken quickly.

The impact was twofold. Engineers were energized by the opportunity to solve a problem of global relevance, and the company gained a system that significantly reduced losses while protecting customer trust. The role of leadership in this context was not to dictate solutions but to create clarity of purpose and provide the latitude for teams to innovate boldly.

Mentorship: A Catalyst for Growth

If leadership provides the framework for innovation, mentorship provides the scaffolding for

Illustration by Stuart Bradford

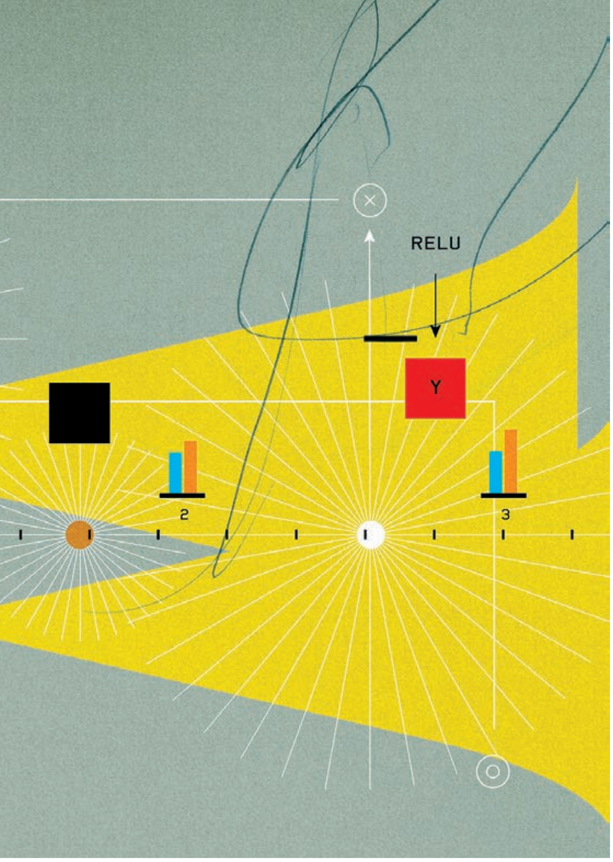

individual growth. Mentorship is not a one-time act but a continuous relationship built on guidance, challenge, and advocacy.

One effective approach I have employed is the use of stretch projects, which are tasks beyond an employee's current skill set, experience, or job responsibilities. For example, when an engineer expressed a strong interest in applying AI to improve our on-call operations, I encouraged him to lead the design and development of a solution leveraging the model context protocol (MCP) standard to reduce on-call workload. MCP standardizes AI models that connect with and use external tools and data sources to automate tasks and simplify integrations.

The effort was successful, attracting contributions from the broader team and reducing the staff's labor for dealing with incidents by more than 1,500 hours annually.

That not only created measurable operational impact but also provided the engineer with a platform to develop his leadership skills and drive innovation at scale.

A feedback-rich environment is advisable. I instituted weekly one-on-one sessions with each of my staff members that extended beyond project updates to cover their career aspirations, strengths, and areas for growth. The conversations uncovered career blind spots, exposed leadership potential, and helped prepare people to step into broader roles.

Mentorship does not stop at giving advice; it involves opening doors to opportunities. I have nominated mentees for conference speaking roles, cross-team leadership positions, and recognition programs.

Encouraging Exploration

Another powerful mechanism to accelerate innovation and growth is intentionally allocating time for self-directed exploration. We dedicated a week every six months during which engineers were encouraged to work on anything they found meaningful, even if it was outside their team or organizational responsibilities.

Some chose to collaborate with colleagues across different departments, while others pursued new projects. The experience gave the employees the freedom to follow their curiosity, sharpen their skills, and explore areas aligned with their personal growth. Beyond skill development, it often led to surprising innovations, as cross-pollination of ideas from different parts of the organization produced creative solutions.

Engineering leadership and mentorship are not optional complements to technical execution; they are fundamental drivers of sustainable success. Leadership provides the vision and structure for innovation, while mentorship nurtures the individuals who bring that vision to life. Together, they create a multiplier effect that advances both technological innovation and career growth. ■

Lokesh Lagudu is an IEEE senior member and a senior engineering manager at Walmart in Sunnyvale, Calif.

LESSONS FOR FUTURE LEADERS

Here are several lessons for those aspiring to lead with impact:

- **Balance technology with people.** Great systems are built by motivated, empowered individuals.
- **Encourage risk-taking within safe boundaries.** Innovation is often the product of bold experiments, not cautious, gradual adjustments.
- **Invest early and consistently in mentorship.** Influence is measured not only by what you build but also by whom you help.
- **Recognize and celebrate achievements.** Acknowledgement fosters motivation, which accelerates innovation and professional development.
- **Create opportunities for exploration.** Allocating time for personal projects can spark creativity and cultivate skills that enrich the organization.

HOW TO THINK LIKE AN ENTREPRENEUR EVERY DAY

Bring startup mojo to your current job

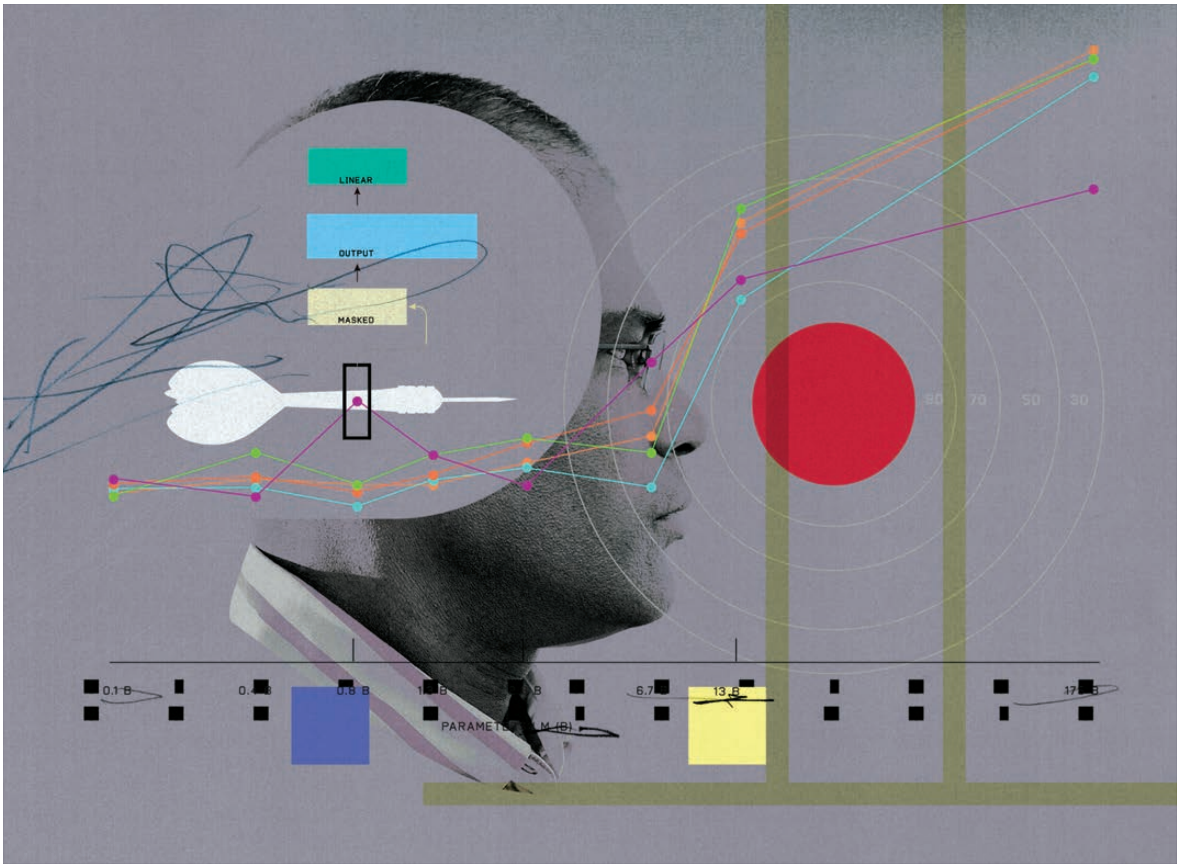

BY ALEXANDER BREM

This article is part of our exclusive career advice series in partnership with the IEEE Technology and Engineering Management Society.

Illustration by Stuart Bradford

LET'S SAY YOU'VE been in your role for a few years now. You know your systems inside and out. You've solved tricky problems, led small teams, and delivered results on time. But lately, between status meetings and routine design reviews, you've caught yourself thinking: There must be a better way to do this task. Someone should make this better.

Then you spend some time imagining. Maybe it's a new tool that would save weeks of engineering time. Or a better process. Or a new product feature. You sketch it out after work hours, maybe even build a quick prototype. Then you think: I could make this product myself.

The shift from "someone should" to "I will" is the start of entrepreneurial thinking. And you don't have to quit your job or have a billionaire's appetite for risk to begin.

Entrepreneurial Thinking

As an engineer, you already have the ability to analyze complex problems, design viable solutions, and follow them through to a working prototype. Your technical skills came from a structured training background and hands-on projects. Your ability to lead, persuade, and navigate uncertainty often comes from experience, especially when you step outside your usual responsibilities.

Some of the most game-changing products didn't begin as formal projects. They started as bootleg efforts—side projects developed quietly by engineers who saw an opportunity. Post-it Notes and Gmail both began that way. Many companies now encourage such efforts; some even allow their engineers to devote 15 to 20 percent of their workweek to pursuing their own ideas.

The Intention-Action Gap

Ideas can be easy. Execution is harder. Nearly every engineer has a colleague with a clever idea that never got past the whiteboard. The difference between wanting to act and actually taking action—known as the intention-action gap—is where entrepreneurship lives or dies. Successful innovators cultivate the discipline to cross the gap, one small, concrete step at a time.

The Innovative Edge

You don't need to be born creative to be entrepreneurial. Here are ways to reprogram your mindset.

> **Challenge the default.** Engineers are taught to follow proven processes, but innovation often starts by asking, "What if we did it differently?"

> **Balance the team.** Innovative companies need a diverse mix of creative thinkers to generate ideas, entrepreneurs to drive execution, and managers to scale efficiently.

> **Know your lane.** Whether you're a visionary, a builder, or an optimizer, understanding your strengths can help you find the right collaborators.

And, yes, timing matters. Amazon might have stayed just an online bookstore without the rise of e-commerce. The right idea at the wrong time is likely to struggle. Start with current trends. For instance, AI offers extremely low entry barriers to get started, and everything is being built around it these days.

An Engineering Attitude

Entrepreneurial thinking isn't only for startup founders. It can mean championing a new process at your company, building an internal tool that changes how your team works, or bringing a product idea from sketch to launch. The engineering mindset—systematic, detail-oriented, problem-solving—is an asset that can power not just products but entire companies.

If you've ever thought, *There's got to be a better way*—and if you felt the itch to make it real—you might be closer to being an entrepreneur than you think. Don't wait any longer; the best time to start is tomorrow. ■

Alexander Brem, an IEEE senior member, is editor in chief of *IEEE Engineering Management Review*, a professor at the University of Stuttgart, in Germany, and an entrepreneur.

IEEE Members Offer Career Advice

They share why IEEE is instrumental and detail keys to their success

> *"Students: Focus on mathematics. Engineering looks hands-on, but math is the foundation behind everything. With practice and persistence, students can succeed and find meaning in the field."*
>
> **—IEEE Member Xiangyi Cheng, assistant professor of mechanical engineering at Loyola Marymount University, in Los Angeles**

> *"Get involved with the key conferences in your technical field of interest. This not only enriches your career but also keeps you connected to the evolving landscape of engineering. Never stop learning about new technologies."*
>
> **—IEEE Senior Member Walt Downing, executive vice president and chief operating officer of Southwest Research Institute, in San Antonio**

> *"Focus on solving real problems, not just building solutions for their own sake. Also build your professional network and seek mentors who can guide you through both technical and career challenges."*
>
> **—IEEE Senior Member Pankaj Gupta, data engineering manager at Discover, in Raleigh, N.C.**

FROM TOP: XIANGYI CHENG; SOUTHWEST RESEARCH INSTITUTE; PANKAJ GUPTA

---

7 Ways New Engineers Can Flourish in the Modern Age

Continued from page 57

5. Develop strong communication skills. Today's engineers don't just code in isolation. You will be expected to:

- Explain design choices to teammates and stakeholders.
- Document decisions clearly.
- Collaborate effectively in cross-functional teams.

This is one area where AI cannot replace you. Clear communication is a career accelerant.

6. Stay curious and keep learning. Cultivate habits:

- Follow industry news, blogs, and open-source projects.
- Experiment with new AI tools, frameworks, and libraries.
- Engage in communities such as GitHub, IEEE Collabratec, LinkedIn, and Medium.

Employers value engineers who keep themselves sharp and relevant.

7. Think beyond coding. AI will increasingly handle routine coding tasks. The differentiators will be:

- Problem-framing: Can you take a vague idea and turn it into a solution?
- Architectural judgment: Can you design systems that scale and last?
- Ethical awareness: Can you spot risks in AI use and address them responsibly?

For more career advice, subscribe to the *IEEE Spectrum* Career Alert e-newsletter at spectrum.ieee.org/newsletters. It features the latest news on jobs, education, and more. ■

Lokesh Lagudu is an IEEE senior member and a senior engineering manager at Walmart in Sunnyvale, Calif.

VERSE BECOMES ELECTRIC **BY PAUL JONES**

NICOLE MILLMAN; SOURCE IMAGE: ISTOCK

Cyborg Laboratory

This is the place where you face yourself,
the you that could be you with a few
different parts, a pump for your heart,
eyes off color, and fresh off the shelf
fake hair (a bit obvious), skin smoothed.
You're not perfect, but it's a good start.

Down to small digits, you'll be improved.
Memory maintained by small motors,
as long as these gizmos don't glitch.
What's before you? Full replacement or
a constant game of test and switch,
pieces peeled off, disconnected, removed,
until you are not yourself, at least,
not the self you knew. That self has ceased,
bit by bit less you at each release.

Paul Jones was inducted into the NC State Computer Science Hall of Fame in 2021 and has published two poetry collections with Redhawk Publications.

澳門大學
UNIVERSIDADE DE MACAU
UNIVERSITY OF MACAU

Full/Associate/Assistant Professor in Ocean Technology

The Department of Ocean Science and Technology (DOST), of the Faculty of Science and Technology, University of Macau, invites applications for the position of Full/Associate/Assistant Professor in Ocean Technology. This position intends to complement and strengthen existing staff expertise. We are seeking candidate with a proven track record in scientific research, an innovative research vision, and a commitment to education.

Ocean Technology (Ref. No.: FST/OST/OT/02/2025)
Position will be open to all areas in Ocean Technology including those areas related to (1) Underwater acoustic communication and network, (2) Ocean remote sensing, (3) Marine imaging and visualization technology, or (4) Marine robotics and autonomous systems.

Applicants must hold a doctoral degree in Electrical and Computer Engineering, Remote Sensing, Ocean Technology, Ocean Science, or related areas. Preference will be given to candidates with research and teaching experience at the tertiary education level.

A taxable annual remuneration starting from MOP1,250,200 (approximately USD154,350) for the position of Full Professor, MOP1,052,800 (approximately USD129,980) for the position of Associate Professor and MOP855,400 (approximately USD105,600) for the position of Assistant Professor will be commensurate with the successful applicants' academic qualification and relevant professional experience. The current local maximum income tax rate is 12% but is effectively around 5% – 7% after various discretionary exemptions.

Applicants should visit https://career.admo.um.edu.mo/ for more details, and apply ONLINE.

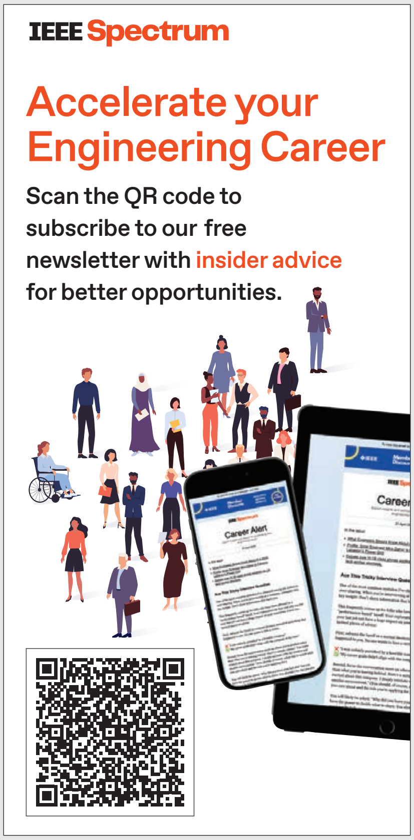

IEEE Spectrum

Accelerate your Engineering Career

Scan the QR code to subscribe to our free newsletter with insider advice for better opportunities.

BY ALLISON MARSH

Past Forward

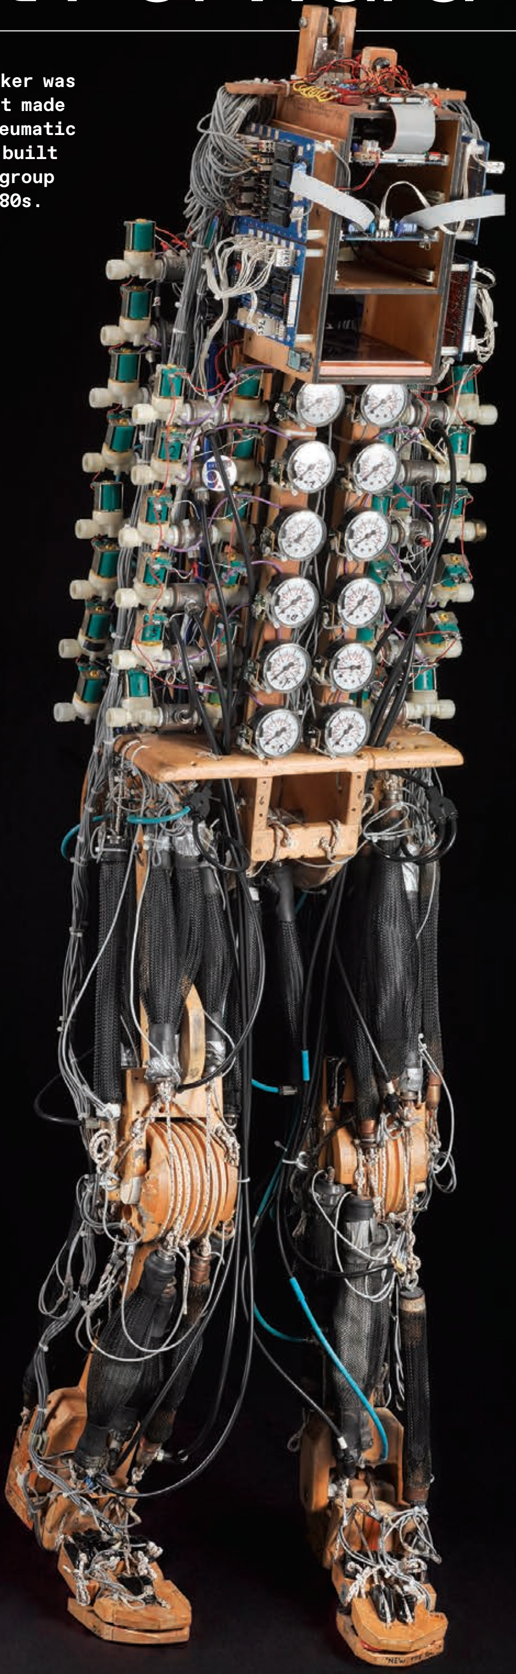

The Shadow Walker was a bipedal robot made of wood and pneumatic "air-muscles" built by a hobbyist group in the late 1980s.

SCIENCE MUSEUM GROUP

Learning to Walk

On a Wednesday night in November 1987, a British photographer named Richard Greenhill invited a group of robot enthusiasts to his attic to tinker with scavenged electronics. Their main objective: to build a two-legged walking robot. Designed by Greenhill's friend David Buckley, a robotics and animatronics expert, the robot used compressed air to extend and contract its 28 "air-muscles." The robot's bones were made of wood, with joints at the hips, knees, ankles, and toes, providing 12 degrees of freedom. The group called itself the Shadow Group, and they soon managed to get the Shadow Walker to stand up reliably and balance itself. But mastering the art of walking proved to be a far bigger challenge. ■

FOR MORE ON THE SHADOW WALKER, SEE spectrum.ieee.org/pastforward-jun2026


THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS
MEDAL OF HONOR


9043
ECHX
CW40-9
USE AH MOTORS ONLY