Spring 2025

Caltech magazine

Kaleidoscope Vision:
Shaping AI at Caltech

Contents

Features

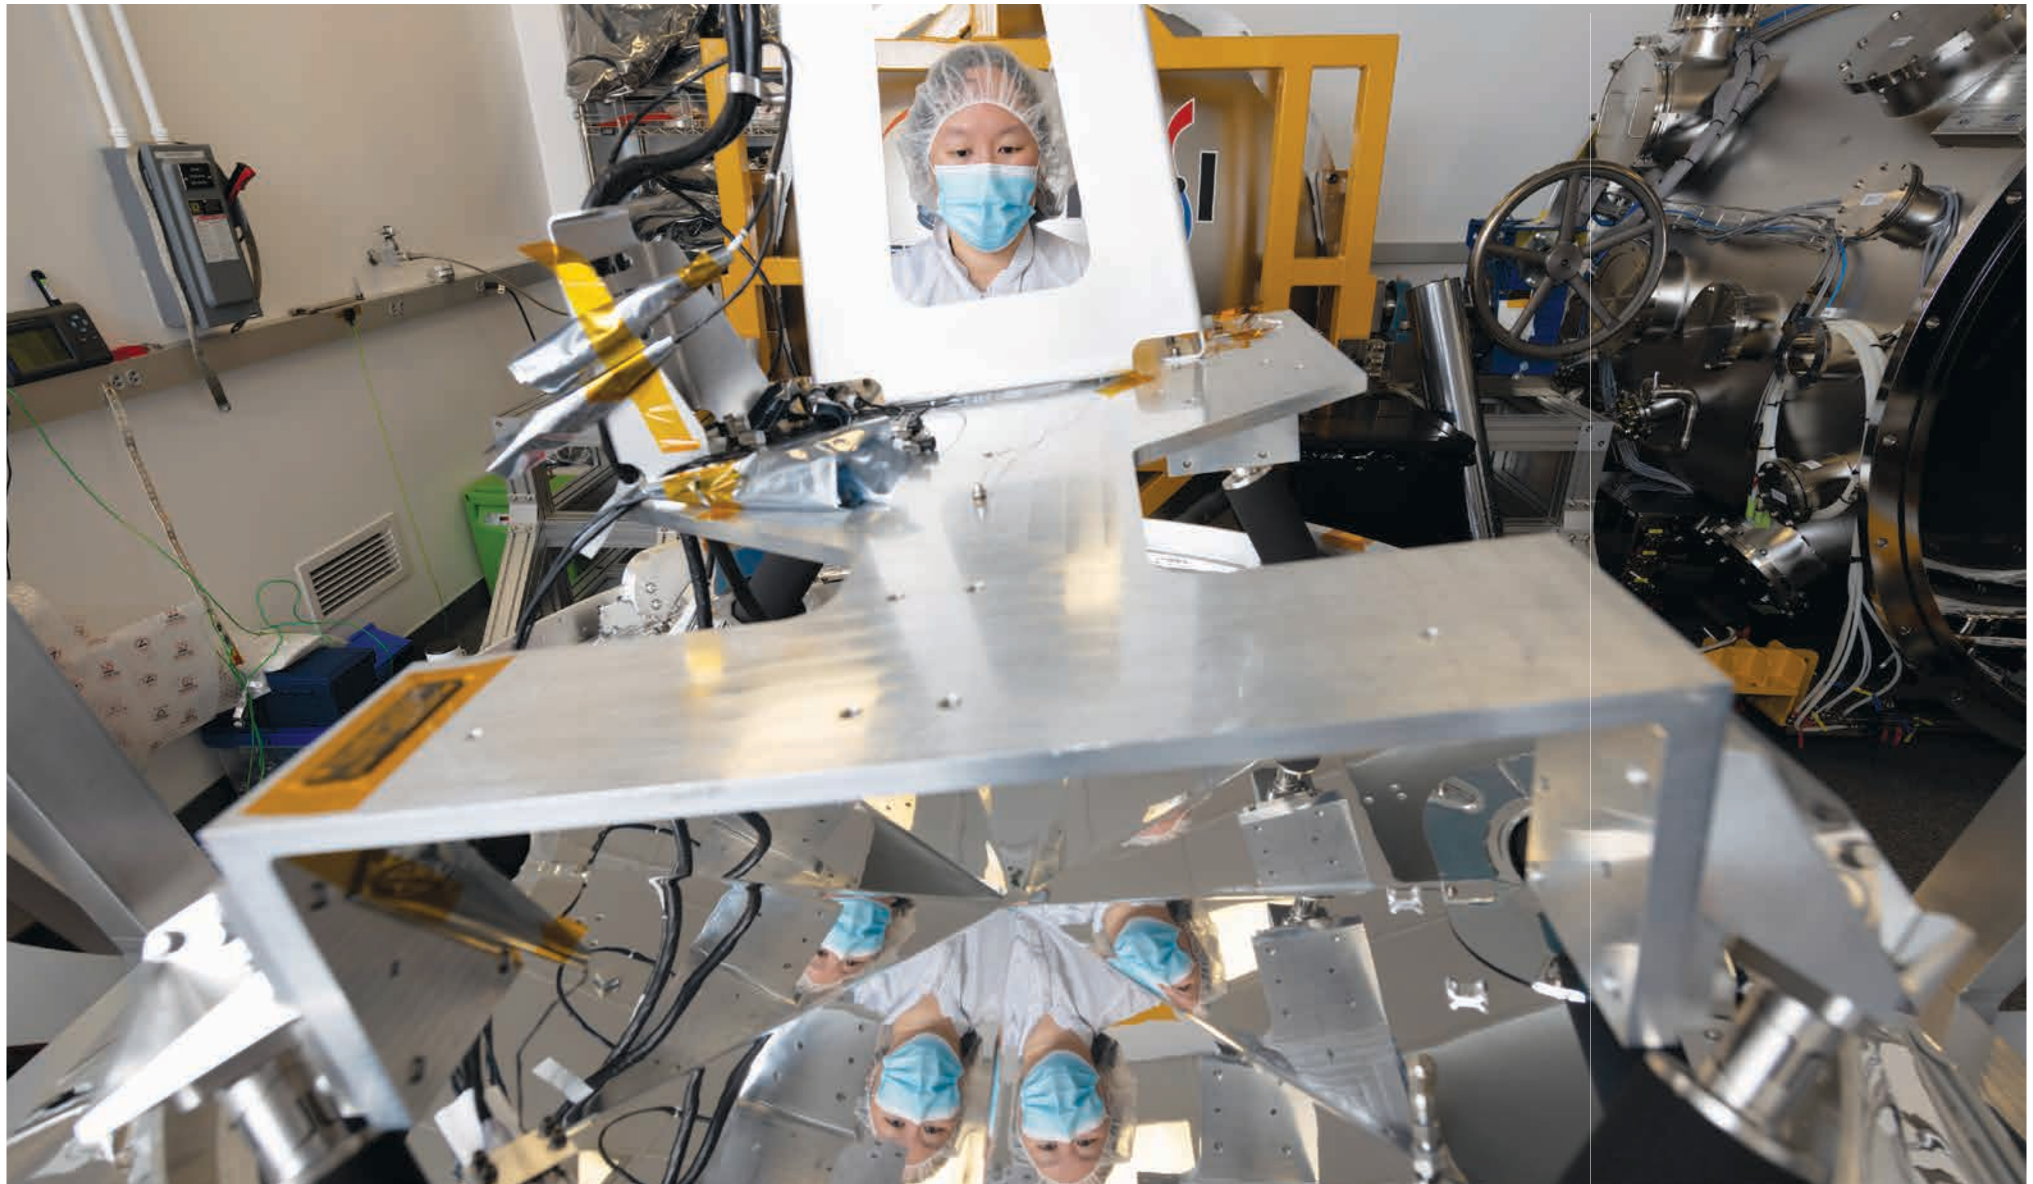

16
Piecing It Together
Looking back at Caltech's prominent role in neural network research that led to the modern AI revolution, and the formation of the Institute's computation and neural systems option.

21
Seeing the Whole Board
Scientists across the Institute are using AI tools in a variety of ways to enable game-changing discoveries.

27
Kaleidoscope Vision
Researchers at Caltech are navigating the ever-shifting ethics of AI in the lab, the classroom, and beyond.

32
A PHOENIX Rises
Caltech lab members installed new air-quality sensors in Altadena after the Eaton fire to monitor particulate levels.

33
Inside Look: Victoria Orphan
Step into the office of this geobiologist, who feels equally at home on campus and miles below the ocean's surface.

36
Behind the Scenes of the Big Bang
Physicist Jamie Bock leads NASA's SPHEREx mission, which has set out to tell us about the beginnings of our universe.

Departments

2 **Letters**

4 **SoCaltech**

12 **In the Community:** Responding to an Unprecedented Disaster

13 **Origins:** How Caltech Launched a Leading AI Conference

39 **In Memoriam**

40 **Endnotes:** How do you think AI will change your job or research field in the future?

Left: Amelia Quan, mechanical integration lead for the Caltech-led SPHEREx mission, is seen with a V-groove radiator, a piece of hardware that will help keep the space telescope cold, at NASA's JPL, which is managed by the Institute, in May 2023.

Online

From animation to ants
Profile: **Meet the Institute's ant wrangler**

Spark of creativity
Feature: **A new kind of Caltech course**

Clearing the air
Article: **Solving an air quality mystery**

Visit **magazine.caltech.edu**

Caltech magazine

EDITOR IN CHIEF
Lori Oliwenstein

SENIOR EDITOR
Omar Shamout

ART DIRECTOR
Jenny K. Somerville

MANAGING EDITOR
Sharon Kaplan

ONLINE EDITOR AND SOCIAL MEDIA COORDINATOR
Andrew Moseman

CONTRIBUTING WRITERS
Whitney Clavin, Lori Dajose, Cynthia Eller, Kimm Fesenmaier, Andrew Moseman, Julia Ehlert Nair, Neil Savage, Omar Shamout

DESIGNER
Jenny K. Somerville

ONLINE DESIGNER
Sergio Solorzano

ILLUSTRATION, PHOTOGRAPHY, VIDEOGRAPHY
Lance Hayashida, Peter Holderness, Jon Nalick, Omar Shamout, Sergio Solorzano, Jenny K. Somerville, Bill Youngblood

COPY EDITORS
Sharon Kaplan, Carolyn Waldron

PRODUCED BY CALTECH'S OFFICE OF COMMUNICATIONS AND EXTERNAL RELATIONS
Shayna Chabner,
Chief Communications Officer

Read *Caltech* magazine on the go at **magazine.caltech.edu**

Contact *Caltech* magazine at **magazine@caltech.edu**

Caltech magazine ISSN 2475-9570 (print)/ ISSN 2475-9589 (online) is published at Caltech, 1200 East California Boulevard, Pasadena, CA 91125.

Image credits: Sergio Solorzano: cover (photo), 8 (upper left), 11; Christian Gralingen: cover, 27 (kaleidoscope design); NASA/JPL-Caltech: TOC (left), 36; Omar Shamout: TOC (upper right); Rob Phillips: TOC (middle right); Shutterstock/trekandshoot: TOC (lower right); Jieyu Zheng: 2; Jon Nalick: 3 (top), 9 (top); Floyd Clark: 3; Matthew Anderson: 4–5; Peter Holderness: 6; Lance Hayashida: 7 (top), 10 (bottom), 12, 22, 24, 29; Wenjie Zhou: 7 (bottom); Julia Gash: 8 (top right); Ketika Garg: 8 (bottom); Iceland Civil Protection/ Björn Oddsson: 9 (bottom); Esther Okamodo: 10 (top); courtesy of Yaser Abu-Mostafa: 13; Jenny K. Somerville: 16, back cover; Bill Youngblood: 17 (top left), 33–35; courtesy Caltech Archives: 17 (lower left, right), 18 (left), 39 (top left, lower left, lower right); Erik Dinel—Allen Institute: 18 (middle); Vicki Chiu: 19; Wikipedia: 20 (top); TED: 20 (bottom); Clément Barbé: 21; Steve Babuljak: 28; Van Urfalian: 30, 31, 39 (top right); SPHEREx/ JPL-Caltech: 37; LindushaJosepha/creative-commons.org/licenses/by-sa/3.0/deed.en: 39 (middle left); Shutterstock/Romolo Tavani: 40 (left); Shutterstock/Roman Samborskyi: 40 (right); Shutterstock/Astrid Gast: 41 (left); Shutterstock/Chokniti-Studio: 41 (right).

Printed by Courier Graphics Corporation, Phoenix, AZ

MIX
Paper | Supporting responsible forestry
FSC® C007548

Pictured on the cover: **Julia Ehlert Nair**

Left: Graduate student **Jieyu Zheng** took this photo of a snowy egret at a Caltech lily pond on January 2, 2025.

Letters

Cycling Celebrity

Readers connected with our #SoCaltech profile of Caltech graduate student and cycling advocate Elina Sendonaris.

Read it here

I'm a JPL-er and Pasadena resident, and just wanted to reach out to thank you for taking that opportunity to promote the Pasadena Complete Streets Coalition [CSC]! I grew up in a car-centric suburb, then spent much of my early adulthood commuting by bicycle in a more walkable city and loving the feeling of freedom and connectedness, and the health benefits it provided. Now I'm living in Pasadena feeling like we have such potential to be a cycling-friendly city but are really missing the necessary infrastructure. I have two young daughters, and I would love to be able to bike with them around town, but for the most part, it just doesn't feel safe.

I signed the petition for the All Ages & Abilities Greenway and will continue to keep an eye on the CSC social media for opportunities to support their work. Thank you for using your platform to promote a worthwhile cause!

–Joseph Russino
PASADENA, CA

Please thank Elina for her efforts to improve cycling. I am doing the same thing in her native Bay Area—Castro Valley. Sadly, it is far more challenging than it should be, and we recently lost a battle to install bike lanes on a spine of our connected bike network because nobody would remove parking—so disheartening. Anyway, I just want her to know there are others out there supporting her efforts.

–Bruce Dughi
CASTRO VALLEY, CA

Laser Focus

It's a Coherent 699 laser in the photograph (not 599). I have used both.

–Mark Lindsay (BS '84)
ESTES PARK, CO

Editor's note: The laser described in our "Shine a Light" story (Fall 2024) was the Coherent 599 dye laser, but the photo wrongly depicted the 699.

Organ Donation

I read George Cannon's letter about Hunter Mead, the Caltech professor who had a large pipe organ in his home [Fall 2024]. As a member (originally) of the class of '76, I was not able to meet him. In fact, the end-of-year Caltech musical event was called the "Hunter Mead Memorial Concert!" I have been the Caltech organist for 50 years. I also served for a while on the LA Chapter of the American Guild of Organists Finance Committee, where I discovered Dr. Mead was one of only two major donors to our fund! Alas, the concerts are now called "Bandorama," and we are forgetting about a fascinating person.

–Les Deutsch (BS '76, PhD '80)
CHATSWORTH, CA

Previous Paleontologists

On page 16 of the Fall 2024 issue of *Caltech* magazine, you state: "Julia Tejada is only the second paleontologist to join the Caltech faculty," the first being Chester Stock. What about Heinz Lowenstam? During my sophomore year (1972–73), I spent time working in his lab. Wasn't he a paleontologist? What I observed of his work would certainly be classified as paleontology.

–Mary Eichbauer (BS '74)
BENICIA, CA

As a former inmate of Jerry Wasserburg's "Lunatic Asylum," I enjoyed the article "Signatures from the Past," but Julia Tejada is not the second paleontologist on the Caltech faculty. She may be the second *vertebrate* paleontologist. As far as I know, Heinz Lowenstam, professor of paleoecology from 1952–83, didn't work on vertebrates, but he was certainly a paleontologist.

–G. Price Russ III (PhD '74)
WALNUT CREEK, CA

Editor's note: Indeed, in "Signatures from the Past" (Fall 2024), we should have identified Julia Tejada as the second vertebrate *paleontologist to join Caltech's faculty. We appreciate those who wrote in defense of Heinz Lowenstam.*

Heinz Lowenstam (at left)

SoCaltech

- Caltech's president preps for transition
- Solid, liquid, or both?
- Superbugs, big personalities, and more

Computer, Start Your Engine

Caltech has a long history of success in out-of-the-ordinary car races. In the Great Electric Car Race of 1968, Caltech's battery-powered Volkswagen bus beat MIT's plug-in Chevy Corvair in a cross-country contest. In 1987, Caltech's Sunraycer won the World Solar Challenge, the first-ever race featuring solar-powered cars.

Now, the Institute can add another achievement to its list. In January 2025, the Caltech Autonomous Systems and Technologies (CAST) Racer made its debut by competing in the Indy Autonomous Challenge (IAC) at the Las Vegas Motor Speedway.

The Caltech racer—an IAC AV-24 IndyCar retrofitted with autonomy hardware and Caltech's control algorithms that allow for autonomous engine/brake control, steering, and navigation—reached a top speed of 155 miles per hour and an average lap speed of more than 144 miles per hour. The roughly 10-person CAST team comprised faculty members, students, postdocs, and staff who notched the impressive high-speed record even though they only formed the full team and started testing in October 2024.

"We are happy with what we achieved because it set a good baseline, so that we can incorporate our advanced autonomy and AI-based software in the next competition. It has taken other teams a few years to reach that speed," says Soon-Jo Chung, the CAST Racer's team leader, Caltech's Bren Professor of Control and Dynamical Systems and a senior research scientist at NASA's JPL, which is managed by Caltech.

Also competing in the race were teams from UC Berkeley, Purdue University, the University of Virginia, Auburn University, Indiana University, Italy's Polytechnic University of Milan and the University of Modena and Reggio Emilia, as well as the Korea Advanced Institute of Science and Technology.

continued on page 6

Computer, Start Your Engine

continued from page 5

Every team in the competition received a functional drive-by-wire IndyCar that they could program for high-speed autonomous driving. After Caltech's IndyCar was delivered in September, the CAST team installed and tested its autonomous driver system, which analyzes data from onboard technology that includes two GPS systems, three LiDAR (Light Detection and Ranging) sensors, inertial measurement units, two radar sensors, and six color cameras.

"With such a high-speed racer, every worst-case scenario for AI-based software is amplified, like road-induced perturbations and turbulence," Chung says. "Everything becomes more challenging. For example, we actually experienced strong winds in Las Vegas, which blew over our heavy monitors and desks. Our racer is best equipped to handle such challenging perturbations because we can utilize our prior machine-learning-based core algorithms to make control decisions a lot faster and safer."

Though the racer was fully autonomous during the competition, autonomous driving is built on human skills, particularly in the areas of data collection and AI software refinement. Caltech staff scientist Matt Anderson, who helped develop the AI software and manage the team, steered the racer during the team's test-drive days on the track. Anderson would ride in the passenger seat of a chase car and assume control of the million-dollar CAST vehicle when it encountered any sudden problems. He did so with a very familiar piece of equipment.

"I drive it with a small Xbox controller," Anderson says. "I've been using one since my parents let me play video games, so I have 30-plus years of experience with it."

The project is also co-led by two other Caltech faculty members, Fred Hadaegh, research professor in aerospace and a senior research scientist and technical fellow at JPL, as well as Mory Gharib, the Hans W. Liepmann Professor of Aeronautics and Medical Engineering, director and Booth-Kresa Leadership Chair of CAST, and director of the Graduate Aerospace Laboratories of the California Institute of Technology. Funding was provided by CAST, Beyond Limits (a Caltech/JPL spinout), and Aramco, among others.

Next up for the CAST Racer is a competition at the WeatherTech Raceway Laguna Seca in Salinas, California, in July 2025.

LA Fires: What Comes Next?

"We know about the problems we face immediately after the disaster. We have to clean out the waste and see what structures can be rebuilt. But we also have to rebuild community, rebuild the businesses. Small businesses are essential to a community. The restaurants and the other types of retail—they are part of who we are and why we feel like we belong in a certain place. Some of them burned. They will try to figure out how to rebuild. But there are a lot of others that weren't burned but are losing their customers—people who can't live here right now, who are busy with recovery, or who are freaked out and start shopping somewhere else. An important piece to a successful recovery is to support those businesses and keep our community alive."

– **Lucy Jones,** a visiting associate in geophysics at Caltech and founder of the Dr. Lucy Jones Center for Science and Society, on community recovery following the LA fires.

Rosenbaum to Retire as President

Thomas F. Rosenbaum, the Sonja and William Davidow Presidential Chair and ninth president of Caltech, will retire on June 30, 2026, after 12 years in the position. He will remain at the Institute as a member of the physics faculty. Under Rosenbaum's leadership, Caltech more than doubled its endowment and completed a historic capital campaign, *Break Through*, that raised $3.4 billion; instituted programs to enhance students' educational and co-curricular experiences and make a Caltech education more affordable; and strengthened NASA's Jet Propulsion Laboratory, which is managed by the Institute.

Rosenbaum ensured that Caltech researchers had the resources and freedom to create, experiment, and define new areas of inquiry, leading to innumerable accomplishments, including the 2015 observation of gravitational waves; the demonstration of the first wireless transmission of power in space, which also led to detectable power being beamed to Earth for the first time; the development of a new type of vaccine that protects against the COVID-19 virus and closely related viruses; and the launch of landmark missions by JPL to study the early history of the universe (SPHEREx) and our closest planetary neighbor (the Mars 2020 Perseverance Rover). During Rosenbaum's tenure, three Caltech faculty members were awarded the Nobel Prize (two of whom were honored in 2017 for the detection of gravitational waves). Institute researchers have also received the National Medal of Science, the Breakthrough Prize, the Kavli Prize, MacArthur Fellowships, and other significant honors.

Rosenbaum also guided the Institute through historic challenges, including the COVID-19 pandemic, a reconsideration of Caltech's admissions practices and the co-curricular student experience, reflection on the Institute's past, public doubts about the value of science, and the devastating Los Angeles fires.

"The path forward has not always been easy or straightforward," Rosenbaum said in a letter to the Caltech community on April 7, 2025, announcing his retirement. "At the same time, these challenges have underscored the value of the skills that we teach to advance societal well-being, educate an involved citizenry, and improve people's lives. … Caltech boasts a solid financial footing, remarkable new research and educational facilities, and, most importantly, the capacity to recruit and sustain the brightest and most innovative individuals from every walk of life and from every corner of the world."

Object Lesson: Polycatenated Architected Materials

Squishy, shear resistant, and solid as a rock? Though it might seem impossible for a material to possess all three of these characteristics at once, that is the case for a new type of matter known as polycatenated architected materials (PAMs). Developed in the lab of Caltech's Chiara Daraio, the G. Bradford Jones Professor of Mechanical Engineering and Applied Physics and a Heritage Medical Research Institute Investigator, PAMs respond to some stresses as a fluid and to others like a solid. They were designed to replicate lattice structures found in crystalline substances, but fixed particles were replaced by entangled rings connected in a stable geometric pattern, similar to chain mail. Because the rings slide against one another when they are not under pressure, they slide into a beaker like a liquid. But when these structures are compressed, they may become fully rigid, behaving like solids. These lattices were 3D printed using a variety of materials, including acrylic polymers, nylon, and metals. Researchers believe PAMs could be used to make helmets and other protective gear as well as biomedical devices and robots.

After the Fires

The Caltech Science Exchange launched a limited-series podcast, *After the Fires,* to highlight what Institute scientists and engineers have learned about the burn zones following the Los Angeles wildfires in January 2025. Since the disaster, researchers from Caltech have deployed sensors to collect information on air quality, tested samples of soil and ash from the burn zone and beyond, monitored debris-flow models that predict areas of concern, studied how the mountains react to rains after fires, and contributed to conversations about community resilience and rebuilding.

Listen to After the Fires for free via Apple Podcasts, Spotify, or your favorite podcast app. For more on Caltech's fire-related research, see page 32.

Visualizing Academic Communities

Research thrives on the exchange of ideas across disciplinary communities, including those at Caltech. Social media can help those groups grow and build new bonds. Ketika Garg, a postdoctoral scholar working on social behavior and social media, wanted to examine these interactions more closely, so she built a tool that visualizes how academic communities connect with one another on BlueSky. The orange nodes in the image below represent BlueSky "starter packs" (lists of accounts related to a topic) that include members of the Caltech community. Connection lines denote how many members the nodes share. "The science we do on a specific topic, in a specific subdiscipline, is part of a much bigger human endeavor with cross-cutting ties and perspectives that make our research richer," Garg says.

Explore the visualization tool

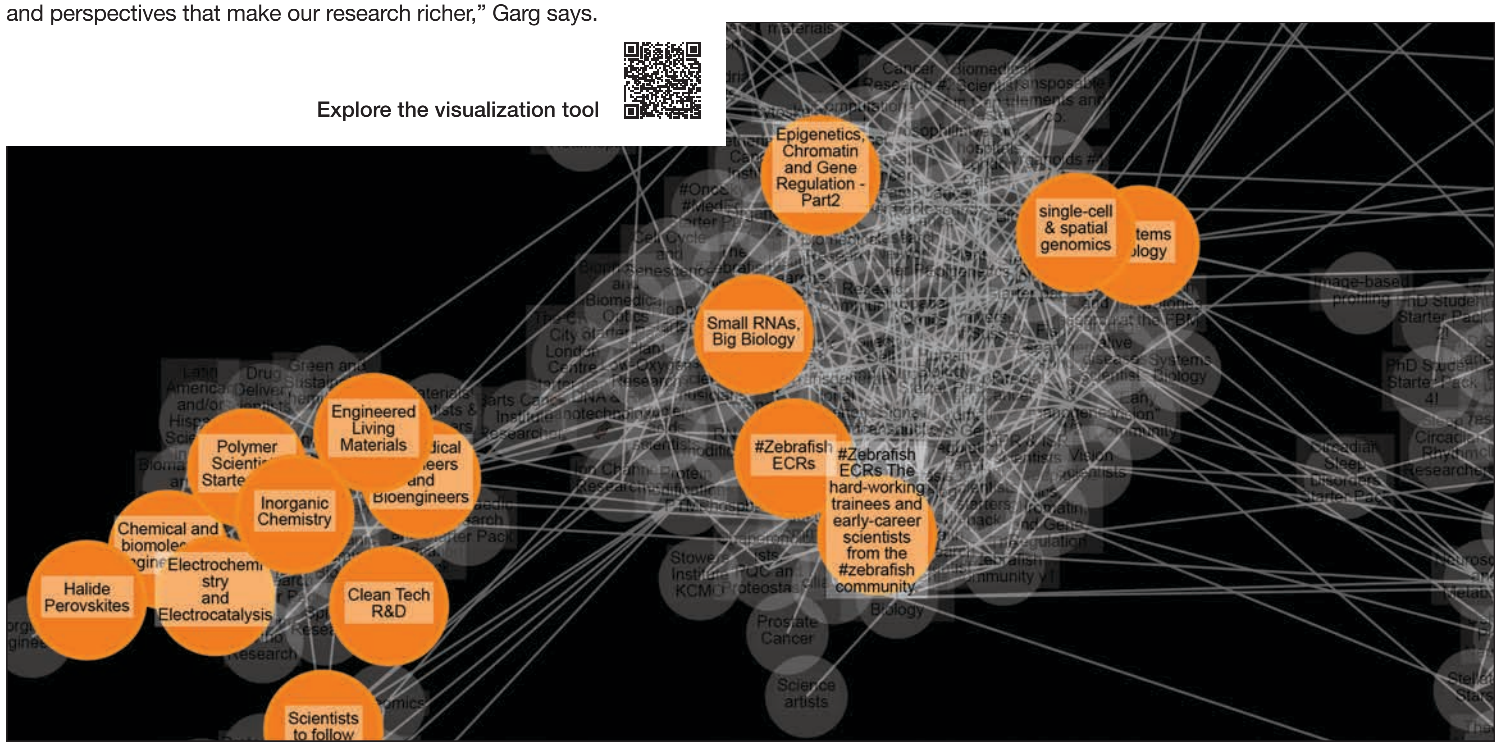

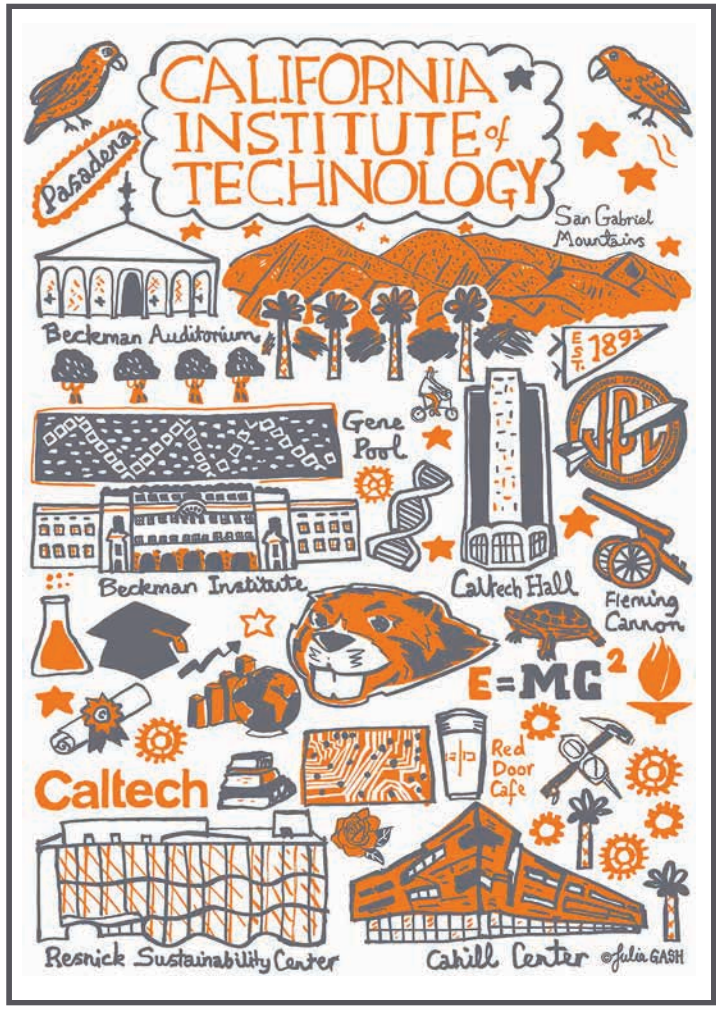

Buy it here

Fire Relief Merchandise

New merchandise featuring a hand-drawn illustration of iconic Caltech and Pasadena symbols by artist Julia Gash can be purchased from the Caltech Store. Fifty percent of the proceeds will go toward the Caltech and JPL Disaster Relief Fund to assist community members impacted by the LA fires.

Early Eruption Warning

The Reykjanes Peninsula at Iceland's southwestern edge is one of the country's most populated regions, and it is also one of the most volcanically active. In 2024, sensing technology developed in Zhongwen Zhan's lab at Caltech was deployed in the region to study the motion of subsurface magma and its eruption into lava on the surface. Using data from the technology, called distributed acoustic sensing (DAS), researchers developed a method to provide warnings up to 30 minutes in advance of lava eruptions. DAS works by pointing lasers into unused underground fiber-optic cables. The image on the right shows an eruption on the peninsula that started on November 25, 2024.

"We run some of the best seismic networks in the world. The Southern California Community Seismic Network (SCSN) runs throughout the region and contains over 500 seismometers. We're planning on adding to networks like these with even more precise tools to measure seismic waves. For example, my laboratory has been developing a method to repurpose telecommunications fibers as a dense array of seismic sensors. With these, we can improve the resolution of measuring earthquakes.

We've also been applying AI to extract information about earthquakes more quickly to better characterize patterns. There is also a lot of work on using both better satellite observations and better computational tools to extract more information from the data."

–**Professor of Geophysics Zhongwen Zhan (PhD '14),** who began a five-year term as the Clarence R. Allen Leadership Chair and director of the Caltech Seismological Laboratory in February, discusses some of what is in store for the lab going forward.

Read more

Which Stealthy Superbug Are You?

Superbugs are pathogens that have acquired resistance genes, making them difficult to treat with antibiotics. Smruthi Karthikeyan, the Gordon and Carol Treweek Assistant Professor of Environmental Science and Engineering and a William H. Hurt Scholar, together with her team, studies these tiny organisms in their natural ecosystems to develop new ways to fight infections. As part of Karthikeyan's Watson Lecture on April 23, 2025, titled "The Secret Life of Superbugs: How Antimicrobial Resistance Moves Between Humans and the Environment," graduate student Grace Solini (below left) created a quiz for guests to find out which superbug best matches their own personality, featuring artwork by her lab mate Sarah Garzione (below right), a grad student in the Caltech–Kaiser Permanente Bernard J. Tyson School of Medicine MD–PhD program.

Take the quiz

Making the Connection: Neurons and AI

The human brain contains approximately 86 billion neurons, which are the primary cells responsible for receiving and sending signals throughout the brain and nervous system. In 1982, John Hopfield, then a professor of biology and chemistry at Caltech, wrote a seminal paper describing how an artificial neural network, modeled after the structure of the brain, could be programmed to learn and to recall. Beginning on **page 15,** find out how the research of Hopfield and others led to the creation of some of today's most popular AI tools, such as ChatGPT, and discover how AI is being studied, engineered, and implemented at the Institute today.

Julian Navarro (fourth-year undergraduate student)

#SoCaltech *is an occasional series celebrating the diverse individuals who give Caltech its spirit of excellence, ambition, and ingenuity. Know someone we should profile? Send nominations to* ***magazine@caltech.edu.***

Navarro is a fourth-year computer science major from San Juan, Puerto Rico. After working as a teaching assistant in 2023 with the First-Year Success Research Institute (FSRI), Navarro co-founded CS Careers @ Caltech, a club that helps students network with recruiters from tech firms and gain professional development skills. After graduation in June 2025, he plans to start a job as a software engineer at Oracle, where he completed an internship in 2024.

"In my first year, I didn't get any summer jobs I interviewed for. I didn't know how to talk. I didn't know what they wanted to hear. Other students had way more coding experience than me coming from high school. When I started teaching at FSRI, I learned how to prepare. I also have a friend at Caltech named Favour Okodogbe who got an internship, and I learned a lot from her about how to communicate in interviews and how to express yourself—like making eye contact and sharing your thought process. That motivated me to help other students who might be in the same position I was in.

Now, with help from Caltech's CALE (Career Achievement, Leadership, and Exploration) and our alumni network, our club brings companies to campus to do workshops. Last year, we brought NVIDIA and Roblox, and we talked about the interview process and the importance of LinkedIn. Then, we set up photo shoots so people can have a good LinkedIn picture. We also made a visit to Google's Venice office with students interested in software engineering and quantum computing. These companies want to recruit here, and we give students the opportunity to connect with them directly in a more personal way."

In the Community

Responding to an Unprecedented Disaster

The devasting wildfires that tore through Los Angeles County in January left more than 300 members of the Institute community without homes, including more than 180 employees at JPL, which the Institute manages for NASA. But in the wake of the fires, people across campus and Lab came together to support their friends and colleagues in any way they could.

Caltech grad students and postdocs started a donation drive that filled up the Hameetman Center with supplies like water and clothing. Many community members welcomed displaced colleagues into their homes—including Caltech President Thomas F. Rosenbaum, who took in then JPL director Laurie Leshin's family when fire threatened their home. At the Institute level, Caltech's Faculty Housing and Student Housing offices helped those displaced by the fire and mandatory evacuation orders to find short- and long-term housing. The Institute also established the Caltech and JPL Disaster Relief Fund to support affected staff, faculty, and students who either lost their homes or needed assistance due to the disaster. As of publication, 1,500 people had applied for assistance, while over 4,000 donors had contributed nearly $5 million. Gifts have ranged from $5 to $500,000, with Caltech alumni contributing more than half a million dollars alone.

Assistance came in other forms as well. Kitty Cahalan, assistant director for educational outreach in the provost's office, left her Pasadena home during the fires to take refuge with her pets and kids at Caltech's Center for Teaching, Learning, and Outreach (CTLO) office. After Cahalan learned that her house in Bungalow Heaven had survived the blaze, she gathered a master list of affected community members and a list of organizations accepting donations.

"I was finding people mostly through social media and text to find out who was affected among our friends," she says. "So, I compiled a spreadsheet of people that we know who were displaced. We were just sending that out to our personal friends, saying, if you want to give, these are the organizations and these are GoFundMe pages of people who we know." When she could return to her own home, Cahalan offered it as a resource. "We said, if anybody ever needs to just come over here, relax, print documents, they should go ahead," she said, noting a friend of her child lived in Cahalan's home for two months until their family could find a more permanent place that could house them all together.

When Mayte Garcia evacuated her Altadena home on the evening of January 7, she did not believe the flames would reach her street. "I live near JPL, and we never thought the fire would get that low," says Garcia, operations manager at the Caltech Center for Inclusion and Diversity (CCID). "It wasn't until a neighbor from our block sent us a video of our house that we realized it did not make it."

After a hotel stay and a week spent living with a family friend, Garcia moved to a new residence in Azusa. During this traumatic period, Garcia's CCID colleagues offered support to her and others whose homes were threatened or lost. "It felt very good that at least one part of my life was stable when everything else wasn't," Garcia says. "They put together a care package and asked what essential items we needed. That felt so good because when your house is gone, you realize you need everything."

Elsewhere, Ralph Adolphs, the Bren Professor of Psychology, Neuroscience, and Biology, organized a fundraiser among his third-floor colleagues in the Chen Neuroscience Research Building to assist custodian Sergio Lopez Meza, who lost his home in the Eaton fire.

"The community has stepped forward in amazing ways," Rosenbaum said at the January campus gathering. "We've been able to open our hearts and open our homes."

–Andrew Moseman

Origins

How Caltech Launched a Leading AI Conference

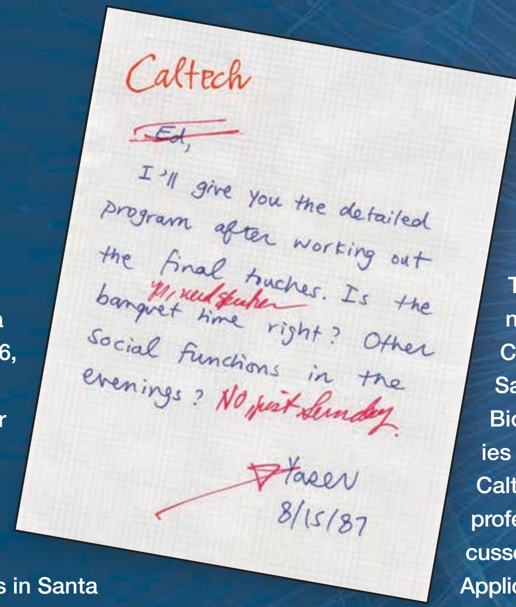

When Yaser Abu-Mostafa (PhD '83) joined the Caltech faculty as an assistant professor after receiving his doctorate, AI had a bad reputation. "Nobody would say they were working in AI because it had promised much and delivered little," says Abu-Mostafa, now a professor of electrical engineering and computer science.

That attitude was about to change, thanks largely to what was happening at Caltech. In 1982, shortly after John Hopfield, now an emeritus professor at Princeton and Caltech's Roscoe G. Dickinson Professor of Chemistry and Biology, Emeritus, published his computer model demonstrating the basic feasibility of artificial neural networks (see page 15), Abu-Mostafa organized a workshop at the Institute to discuss the model and its implications.

"If you were a scientist in this area, and you looked at the whole picture, you said, 'This will lead to something. I don't know what, but it'll lead to something, and we'd better pursue it.'"

–Yaser Abu-Mostafa on the AI research presented at the first NeurIPS conference in 1987

Today, that event with just 75 attendees has grown into what is called the Conference of Neural Information Processing Systems (NeurIPS), the leading gathering of artificial intelligence researchers and professionals in the world, which attracts thousands of attendees every year.

"The Hopfield model was very inspirational for us," says Abu-Mostafa about the recent Nobel laureate's work. "We wanted to analyze the capabilities, see what could be done with it, and so on." After two more years of small gatherings, the workshop had a growth spurt. In 1986, Abu-Mostafa's former Caltech advisor Demetri Psaltis, together with Bell Labs, co-hosted a larger meeting of about 150 scientists in Santa Barbara. "That's when we realized this is serious business," Abu-Mostafa recalls.

In November 1987, Abu-Mostafa and Ed Posner, the late Caltech professor and JPL chief technologist, organized the first NeurIPS conference in Denver, Colorado, with sponsorship from the Institute of Electrical and Electronics Engineers Information Theory Society. Posner served as the founding general chair. As program chair, Abu-Mostafa says he "sweated bullets" knowing the conference's success hinged on the quality of the research presented. He was so selective that he nearly rejected a paper submitted by a member of the conference organizing team and only permitted himself to present a poster about the relationship between entropy and connectivity in neural networks. "I invested a ridiculous amount of time in the program, and it paid off," he says.

Around 600 scientists attended the first official conference, including computer scientists, biologists, mathematicians, and engineers. The first two plenary speakers were Carver Mead, Caltech's Gordon and Betty Moore Professor of Engineering and Applied Science, Emeritus, who spoke about the engineering side of "Networks for Real-Time Sensory Processing," and Terry Sejnowski, now the Francis Crick Chair at the Salk Institute for Biological Studies and a former Caltech visiting professor, who discussed "Biological Applications of Neural Network Models."

"There was enough substance that if you were a scientist in this area, and you looked at the whole picture, you said, 'This will lead to something. I don't know what, but it'll lead to something, and we'd better pursue it,'" Abu-Mostafa says.

The conference remains a primary forum for scientists and engineers to present and discuss developments in AI. Now organized by the NeurIPS Foundation, established by Posner, the 38th conference last December in Vancouver, British Columbia, saw more than 16,000 registrants attend along with nearly 3,000 others online.

"Every breakthrough over the last 37 years was presented at NeurIPS," Abu-Mostafa says. These include AlexNet (2012), a model developed by a University of Toronto grad student that roundly outperformed all computer vision programs of the time; AlphaGo (2019), the Google DeepMind program that defeated human champions of the Chinese game *Go*; and the model for ChatGPT (2022).

"AI is no longer a bad word," Abu-Mostafa says. "It is either an exciting word or a scary word. It is incumbent upon us, the AI researchers, to make sure that the exciting part flourishes and the scary part is eliminated or at least mitigated."

–Kimm Fesenmaier

The Many Facets of AI

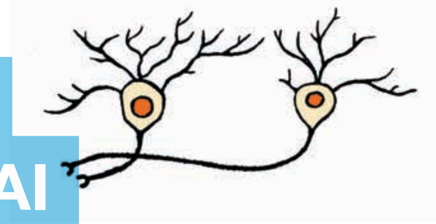

When you look into a kaleidoscope, play a game of chess, or open an Erector Set, you are unlocking new perspectives and possibilities. Each twist of a hand, move on a board, and connected piece represents a single step in a journey that could produce any number of different outcomes. In much the same way, artificial intelligence draws on and reconfigures data, **generating seemingly limitless insights** while simultaneously **mirroring the biases** we put into the system. No longer a mere curiosity, AI has emerged as a **foundational tool of the educational, research, and societal landscape.** Like the rest of society, Caltech's scientists, educators, and administrators are navigating this new reality. But they are not passive observers; they are **pioneering efforts to harness the technology in innovative ways** as they seek to **solve fundamental scientific mysteries** and **forge new lines of inquiry.** Meanwhile, they are **weighing the implications** of AI's increasingly ubiquitous presence. The three features that follow seek to illuminate this evolving journey on campus, while also tracing how the Institute helped enable this paradigm shift.

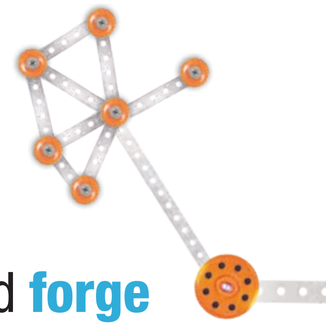

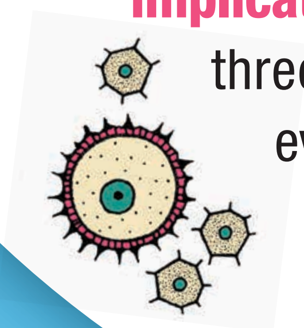

Piecing It Together

Tracing the roots of neural networks, the building blocks of modern AI, at Caltech.

By Whitney Clavin

In the early 1980s, three giants of Caltech's faculty—Carver Mead (BS '56, PhD '60), now the Gordon and Betty Moore Professor of Engineering and Applied Science, Emeritus; the late Nobel Laureate in Physics Richard Feynman; and John Hopfield, then a professor of biology and chemistry, who would go on to win a Nobel Prize as well—became intrigued with the connections between brains and computers. The trio would gather for lunch at Caltech's Athenaeum and wonder: How do our brains, with their billions of interconnected neurons, process information? And can computers, which work in more straightforward number-crunching ways, mimic the brain's ability to, in essence, think?

These conversations ultimately led to a new graduate-level course, The Physics of Computation, intermittently co-taught by all three professors from 1981 to 1983. Hopfield recalls that there were obstacles to getting the course off the ground. "Back then, there was very little interaction between computer science and other fields," says Hopfield, now an emeritus professor at Princeton and Caltech's Roscoe G. Dickinson Professor of Chemistry and Biology, Emeritus. "We had a diffuse mixture of ideas we wanted to present, and it took a while for us to convince the Institute to approve the course. Still, it was an intellectually exciting period and brought great new students and many guest lecturers to Caltech."

Alongside these conversations, Hopfield began to formulate ideas for creating simple networks that mirror the way human memory works. In 1982, he published a theoretical paper describing how an artificial neural network, modeled after the structure of the human brain, could be programmed to learn and to recall. Though other researchers would later build these networks using computer chips, Hopfield's research used math to describe a new biology-inspired scheme that could be trained to "remember" stored patterns, such as images. The computers could recall the images even when only an incomplete or fuzzy version of the same image was available. The system is akin to someone remembering the full experience of

Clockwise from top: Carver Mead (BS '56, PhD '60), **John Hopfield,** and **Richard Feynman.**

hearing a particular song after catching a snippet of the tune on the radio.

The roots of modern AI programs like ChatGPT can be traced to biology-inspired models similar to the Hopfield network, as it is now known. For this seminal research, Hopfield was awarded the 2024 Nobel Prize in Physics, together with Geoffrey Hinton of the University of Toronto.

Hopfield's breakthrough came during a pivotal time in Caltech's history when ideas had just begun to flow between neuroscience and computer science. "AI research was developing very slowly and still had many doubters," Hopfield says. Despite the challenges, Hopfield sought to make the movement more official:

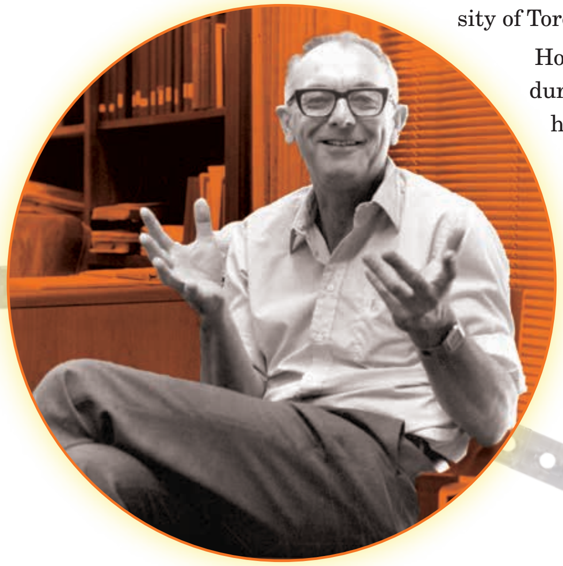

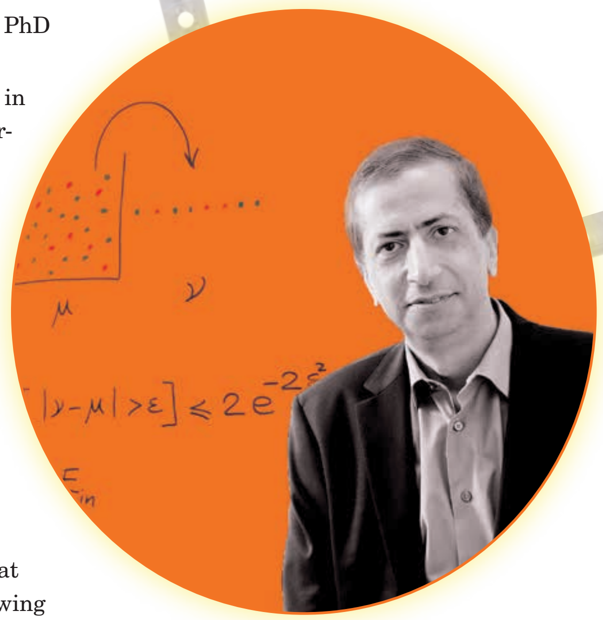

From left to right: **Robbie Vogt, Christof Koch, Yaser Abu-Mostafa** (PhD '83), and **Erik Winfree** (PhD '98).

In addition to the Physics of Computation course he helped run, he sought to organize a new interdisciplinary program offering graduate-level degrees.

Caltech's then-provost Robbie Vogt, now the R. Stanton Avery Distinguished Service Professor and Professor of Physics, Emeritus, supported the idea, and, in 1986, the Institute's Computation and Neural Systems (CNS) program was born, with Hopfield as its first chair. Today, CNS comprises a vibrant group of scholars that has produced more than 150 PhD graduates.

"This program was the first of its kind that took in highly quantitative students from physics, engineering, and mathematics who were interested in both brains and computers," says Christof Koch, who served as the first faculty member hired in CNS and later on as a chair of CNS, before leaving Caltech in 2013 to become the chief scientific officer and president of the Allen Institute for Brain Science. "Now there are many other places that similarly look at brains as computational systems, but we spearheaded the effort."

Yaser Abu-Mostafa (PhD '83), a professor of electrical engineering and computer science at Caltech who did theoretical work on Hopfield networks in the 1980s, recalls that by the middle of that decade, more and more people were joining the growing AI community worldwide thanks to the innovative work being done on campus. "What Hopfield did was very inspirational," he says. "It established in people's minds that this can be done." Abu-Mostafa initiated an AI-themed workshop, which later led to the creation of the Neural Information Processing Systems conference in 1987. Now known as NeurIPS, the gathering has become the largest AI conference in the world. (See page 13.) "It has been very rewarding to watch a field forming from scratch," Abu-Mostafa says.

Built on Physics

In the late 1970s, Hopfield, then a biophysics professor at Princeton University, attended a series of neuroscience lectures in Boston and quickly became fascinated with the topic. As a condensed matter physicist by training and the son of two physicist parents, he wanted to understand how our minds emerge from the complex network of neurons that make up human brains. "I was very interested in the interface of physics and living matter," he says.

In 1980, Hopfield left Princeton for Caltech in part due to the Institute's "splendid computing facilities," which he would use to test and develop his ideas for neural networks. However, Hopfield did not set out to create an artificial intelligence. "I was hoping the networks would tell us how the brain works," he says.

His idea was to build a simple computer program based on the vast network of billions of neurons in our brain and the trillions of connections among them. Computers of the 1980s were used to execute long sequences of commands and search databases for information, but that process took time and required increasingly large amounts of storage space. Imagine trying to remember the name of a singer and having to comb through a catalog of all the singer names in your head one by one—it might take a while.

Instead, our brain has a more efficient system of retrieving information that relies on neurons changing their architecture as they learn new connections. Memories are encoded in different patterns of neural activity; as Hopfield says, the brain is a dynamic biological system.

He decided to model his neural network on another dynamic system in nature involving magnetism. Called the Ising model, the system describes how the up or down spins of electrons in a material can influence one another and spread magnetized states. When this occurs, the system evolves toward the lowest-energy state, like a ball rolling down a hill.

Hopfield networks also evolve toward low-energy states in a mathematical sense. These neural networks are composed of artificial neurons connected via nodes, with each connection having a different strength, or weight. A set of computer codes, known as an algorithm, directs the network to tune the connection strengths between these neurons such that a stored image, like that of a spider, becomes linked to a particular low-energy state. When a fuzzy image of a spider is fed into the Hopfield network, the network's artificial neurons assess the available information and then adjust their activity levels by evolving toward the low-energy state matching the stored image. In this way, the system learns to recognize images of objects.

The backbone of any neural network is an algorithm (or learning rule); a key feature of Hopfield's algorithm, says Abu-Mostafa, is that it allowed the system to learn and grow increasingly smart. "Learning is absolutely essential to intelligence," he says. "Hopfield extracted the essence of neurons." Abu-Mostafa notes that the theoretical paper Hopfield published in 1982, "Neural networks and physical systems with emergent collective computational abilities," is the fifth-most-cited Caltech paper of all time.

Physics played a key role in Hopfield's success, Koch says, and this "led to a massive influx of physicists into the field."

"Hopfield figured out how to mold the energy landscape [a map of the possible energy states of a system]. His network was trained to dig a hole in the landscape corresponding to the image pattern being trained," adds Erik Winfree (PhD '98), professor of computer science, computation and neural systems, and bioengineering at Caltech, and a former CNS student of Hopfield's. "He brought physics to the networks."

In Hopfield's Nobel Prize lecture in December 2024, he explained how the Ising model of magnetism could be generalized to replicate a biological system like the brain. "Everything really came together when I saw these two parts of science are really described by the same set of mathematics," Hopfield said.

Mead adds that others had attempted to build artificial neural networks before, but few could envision them scaling up to the sizes needed to perform interesting tasks. "Hopfield showed that they were possible," he explains. "This was the first time people started thinking that neural networks might be useful."

How Computers Caught Up

Around the time that Hopfield was working on the theory behind his neural networks, Mead and his collaborators had begun to transform the computer industry by inventing a new way to pack more of the tiny semiconductors known as transistors onto computer chips, a process called very large-scale integration (VLSI). VLSI allowed millions, and now billions, of transistors to squeeze onto single chips, a feat that enabled the development of desktop computers, cell phones, and myriad other computing gadgets.

In the early 2010s, researchers realized they could use a type of VLSI chip employed in video games, called graphics processing units (GPUs), to handle the huge computational demands of AI networks. Though GPU chips were not invented at Caltech, some aspects of their origin can be traced back to the early VLSI research on campus.

A key feature of GPUs, which makes them critical for large AI neural networks, is a type of computing called parallel processing. Essentially, this means they can

From top: **Bill Dally** (PhD '86) and **Anima Anandkumar.**

perform multiple computations at the same time, making them very effective for solving math problems. This innovation came from a computer scientist working with VLSI technology in the 1980s named H.T. Kung. Then a faculty member at Carnegie Mellon University and now at Harvard University, Kung gave a talk at the first VLSI conference.

"He figured out how to multiply whole rows of numbers, not just two at a time, on the VLSI chips," Mead explains. "It's called matrix multiplication, and it allowed for parallel processing. The idea was later rediscovered by NVIDIA and turned into GPUs."

NVIDIA, the world's leading developer of GPUs, also has its share of Caltech influences, including Bill Dally (PhD '86), a former Caltech professor who is now the company's chief scientist and senior vice president, and Anima Anandkumar, Bren Professor of Computing and Mathematical Sciences, who previously served as the company's senior director of AI research from 2018 to 2023.

Like Hopfield, Anandkumar says physics inspires her work. Even before Anandkumar joined Caltech in 2017, she says she "was fascinated by physics." In 2011, she analyzed how the success of learning algorithms is tied to the phase transition in the Ising model, the same model upon which Hopfield built his network. "Hopfield gave us the starting tools for modern AI," Anandkumar says.

Building Bridges Between Brains and Computers

Hopfield points to Mead as an early believer in his vision for neural networks. "Carver had me give a talk in the 1980s where people from Bell Labs would be," Hopfield says, "and I remember thinking, I don't know what to tell these people. Then I realized I could simply prove the theorem for the Hopfield network. The original proof is written on the back of hotel stationery that I still have."

Vogt, the Caltech provost during this time, also believed in the viability of Hopfield's efforts and ultimately green-lighted the formation of the CNS graduate option. "I don't think CNS would have gotten going for another year or two if it hadn't been for Robbie Vogt," Hopfield says. "He was a different kind of leader. He could do marvelous things."

Hopfield saw CNS as a means for people with different backgrounds to converse and influence one another's work, though he notes it was difficult to get both the Physics of Computation course and the CNS graduate option launched at Caltech. Other scientists, he says, were not convinced of the merits of the interdisciplinary effort.

"Before CNS, there was a clear gap between computer science and neurobiology," he says. "The gap was something like having a set of people working on weather, and another set of people working on molecular physics and chemistry but having no one asking what the relationship was between weather and the molecular collisions, which were obviously at the bottom of it. The quality of the CNS incoming students was so high that the neurobiologists and engineers who had been skeptics rapidly became true believers, or at least willing participants."

Today, the nearly 40 faculty members associated with Caltech's CNS graduate option continue to study the human brain as a computational system both to develop new AI tools and better understand the brain's own fundamental workings.

At a 30th anniversary celebration for the program held in 2017, many graduates of the program remembered the excitement of crossing boundaries between fields. Gabriel Kreiman (PhD '02), a professor at Harvard Medical School, spoke at the event and credited the program's rigor and collaborative nature for producing great science and scientists.

"The intellectual freedom to get together and go with all the other CNS people to the Athenaeum to have lunch and then spend three hours discussing the minutiae of one particular problem, or staying until the wee hours in one of the rooms where we have all of the computers and working together and fighting together about absolutely every problem in neuroscience and computational neuroscience…" said Kreiman at the event, "the magic, the spark of what happened here at CNS was completely unique." [C]

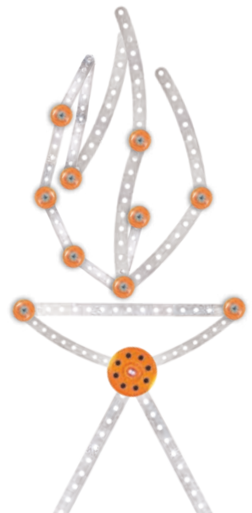

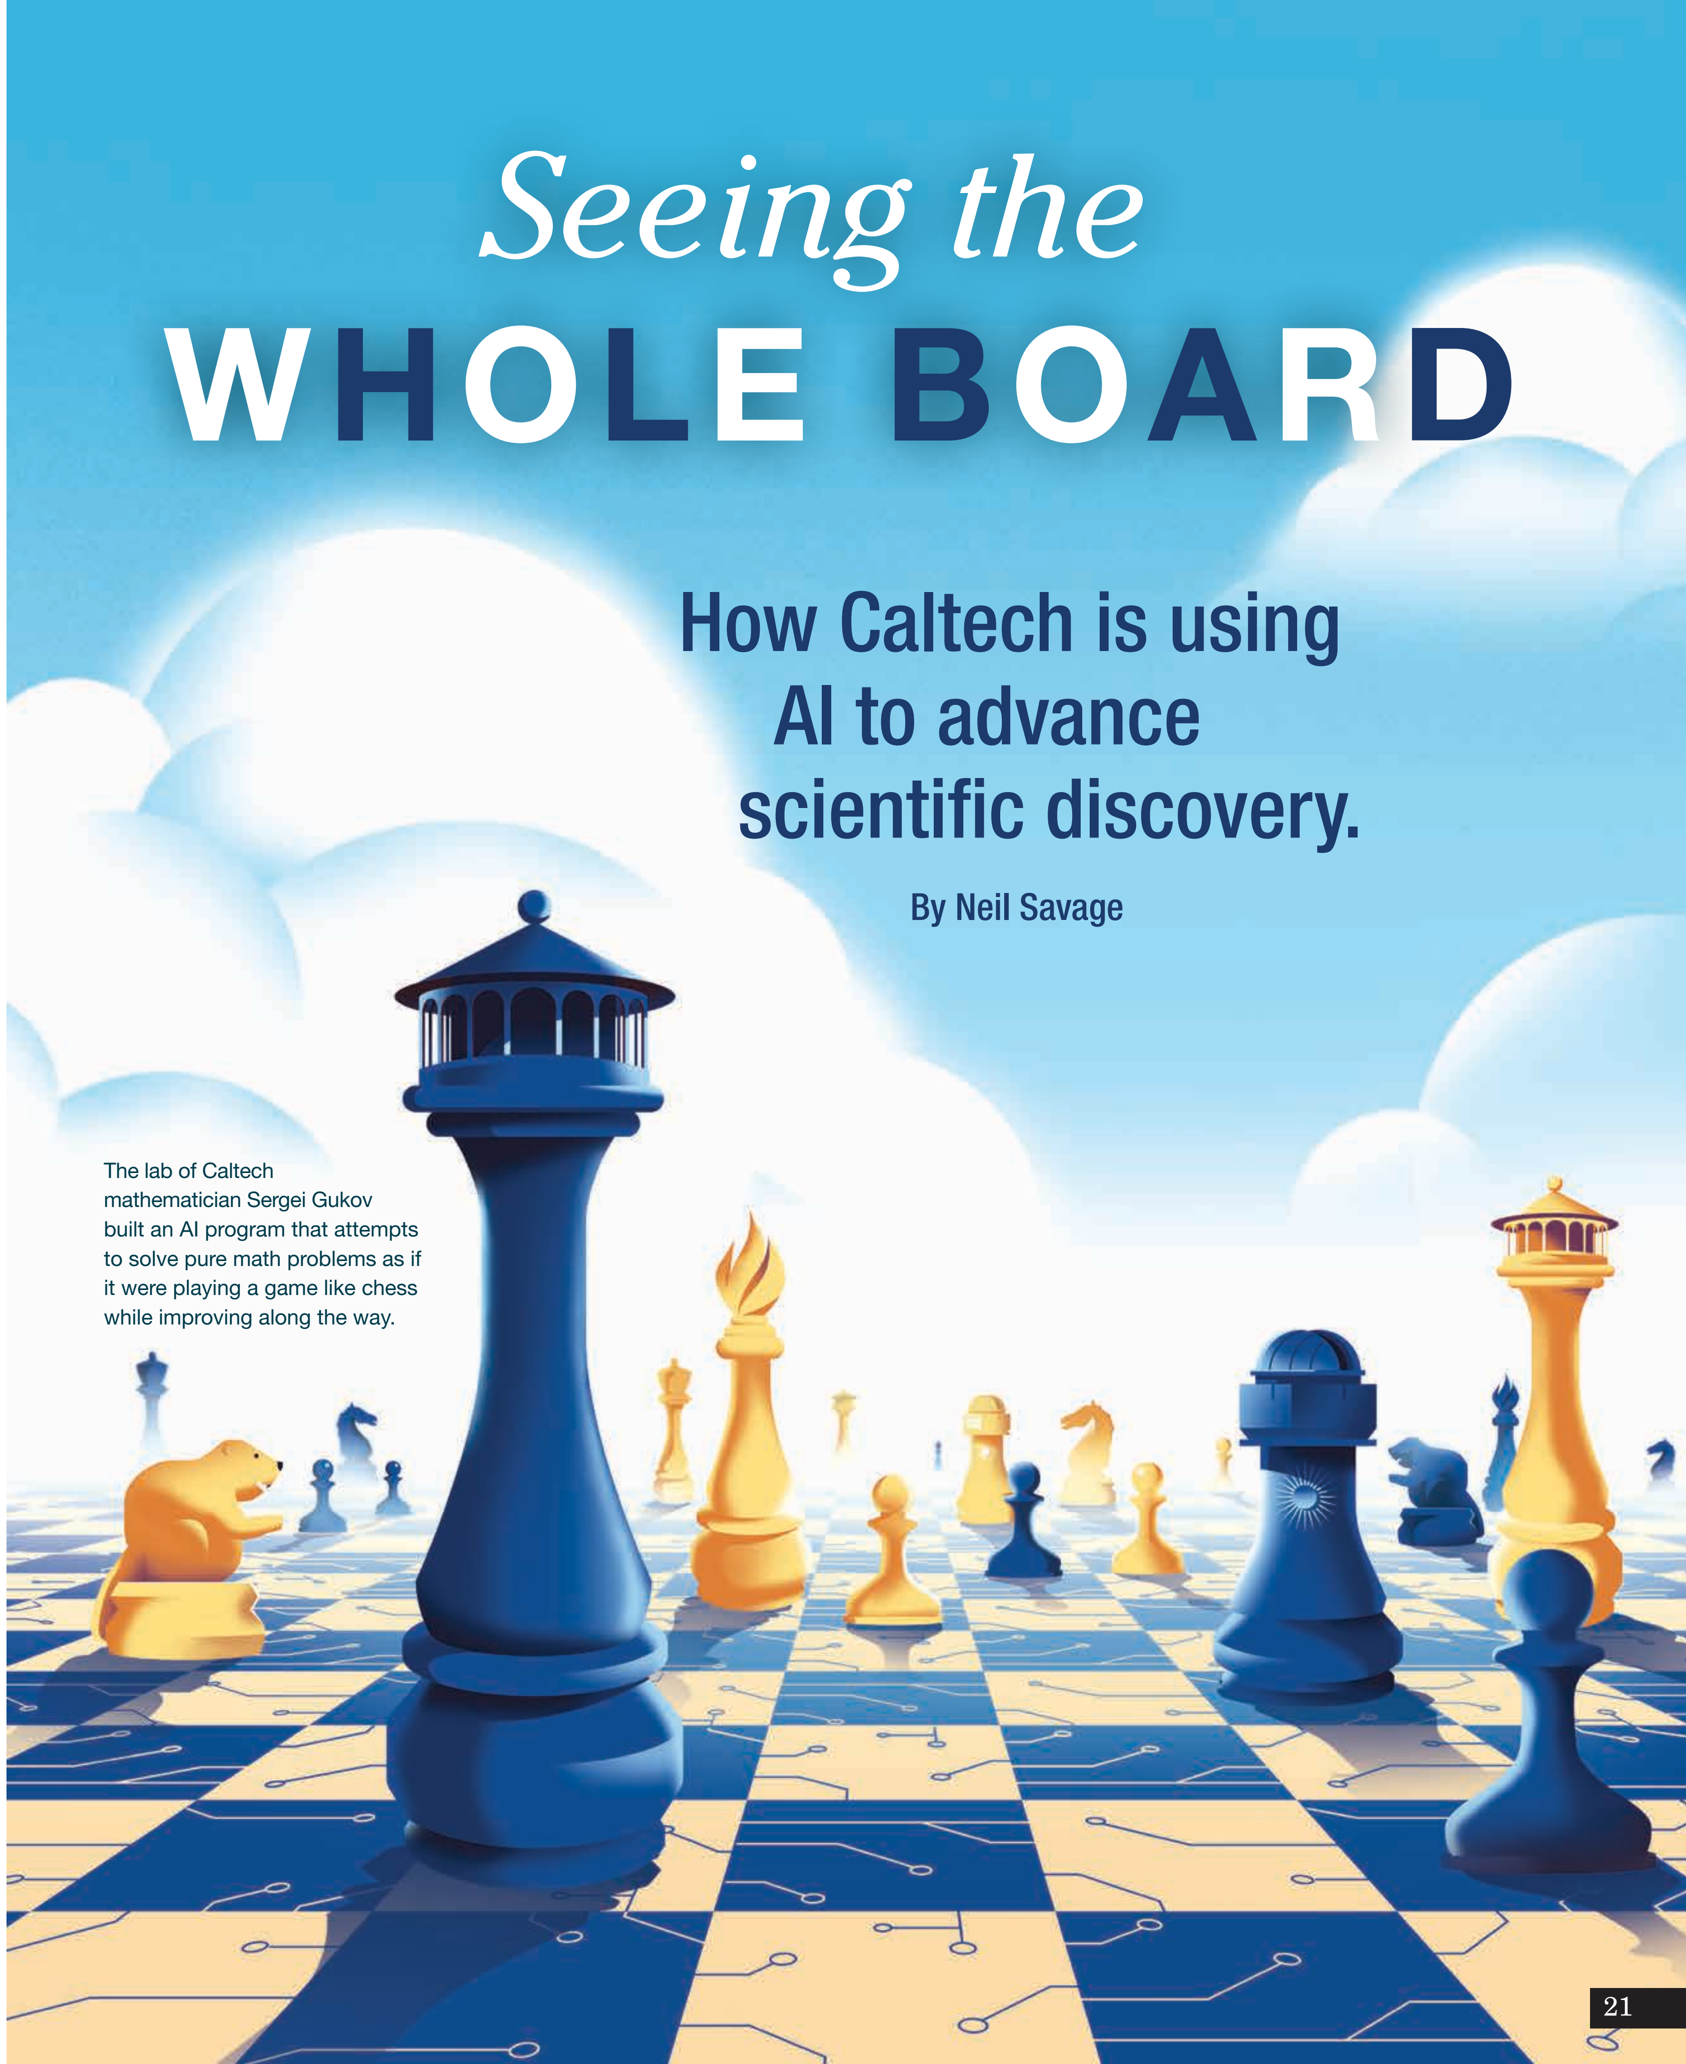

Seeing the WHOLE BOARD

How Caltech is using AI to advance scientific discovery.

By Neil Savage

The lab of Caltech mathematician Sergei Gukov built an AI program that attempts to solve pure math problems as if it were playing a game like chess while improving along the way.

Some five years ago, mathematician Sergei Gukov began teaching himself how to build the neural networks that are the foundation of artificial intelligence, simply to see whether they might be useful in the realm of pure mathematics. He was, he admits now, skeptical whether the supremely complicated and complex questions being asked by pure math would be within the reach of AI's ability to process information.

Today, he says, he is no longer a skeptic, thanks to the time he has spent learning how to build those neural networks, in what he believed was an effort to show that they were irrelevant to the kind of work he was involved in. "When we think about people working on pure math, we usually think of someone sitting in the attic and proving theorems that are so esoteric that no other human can understand them, let alone machines," says Gukov, Caltech's John D. MacArthur Professor of Theoretical Physics and Mathematics.

Instead, he proved to himself that the opposite is true. To understand why, Gukov says it is important to recognize that solving hard math problems can be thought of as a sort of game. These problems involve assertions that mathematicians believe should be true, and their challenge is to prove that they *are* true. In other words, all these problems are essentially a search for the path from A to B. "We know the hypothesis, we know the goal, but connecting them is what's missing," he says.

What makes these problems so hard is the number of steps from A to B. Whereas an average game of chess lasts about 30 to 40 moves, these problems require solutions that take a million or more steps, or moves. After studying neural networks, Gukov realized he could build an AI algorithm that learns to play the game better, or solve a particular problem, as it competes against itself. The program starts with knowledge of its existing conditions, a set of rules about what moves it can make and a definition of what it means to win. Then it uses a machine-learning technique called reinforcement learning, similar to the way a person might train a dog to sit, in which the computer tries a move and gets feedback on whether it is closer to its goal.

Gukov and his colleagues recently used this approach on a decades-old problem known as the Andrews-Curtis conjecture. They did not actually solve that problem, but they managed to disprove two sets of potential counterexamples that, had they been true, would have disproved Andrews-Curtis. Though many of these math problems seem intractable now, Gukov says that in 10 years finding those paths from A to B could seem as simple as a computer winning a game of chess, a feat once considered nearly impossible. "I don't know how likely or unlikely it is, but there is definitely a chance that AI can master that kind of problem," he says. "It develops new strategies that are better than what humans can do."

Across the Institute, researchers are learning that AI can help them to think bigger and do more. They are confident about the varied roles the technology can play in scientific research, whether that involves crystallizing mountains of data into new and useful insights, uncovering patterns in data too subtle for humans to notice, or using the power of neural networks or machine learning to streamline experimentation and create new knowledge, develop new therapies, or understand the world and its complex systems in new ways.

Anima Anandkumar, Caltech's Bren Professor of Computing and Mathematical Sciences, says Caltech has taken a leading role in imagining and implementing new uses for AI. "I've been at Caltech since 2017, and since then, I've had so many collaborations, including with other faculty who work on AI, and interdisciplinary collaborations across campus," she says. "The small size of Caltech and its open mindset has made this possible much more than in other, bigger universities. The impact of it can be directly seen in the work that has happened since."

The Institute has infrastructure in place to assist labs as they make this transition. Anandkumar, along with Yisong Yue, professor of computing and mathematical sciences, run AI4Science, an initiative launched in 2018 that helps scientists across the Institute discover which AI tools and resources might help advance their work. The program is a collaborative effort with the University of Chicago and has support from the Margot and Tom Pritzker Foundation. Seven years on, Anandkumar and Yue have partnered with principal investigators from a wide variety of disciplines on a dozen projects in which AI has helped to make significant advances. "Fundamentally, AI is already transforming the whole scientific method," Anandkumar says.

As part of AI4Science, Yue teaches a class that provides PhD students with a basic understanding of how AI works and what kinds of AI tools they can incorporate into their research. One point he stresses is the need to start with high-quality data. "Data is the fuel for AI," Yue says. "AI converts that data into this model from which you can extract knowledge." If researchers lack good data, he helps them find it, whether that entails digging it out of the scientific literature, using a computer to generate simulated data, or more effectively leveraging the collected experimental measurements.

Seismologist **Zachary Ross** (left) works with computer scientist **Yisong Yue** to develop AI visualizations of underground seismic activity.

AI to Model the Physical World

The wildfires that swept through the Los Angeles area in January, touched off by an extreme wind event, reinforced the importance of timelier weather forecasts. Anandkumar and her colleagues are working on AI technologies to create those forecasts and potentially save lives in future natural disasters.

Existing weather-forecasting models are based on complex mathematical equations that describe the physical processes governing how Earth's atmosphere and oceans behave. They are fed by observations of current conditions, such as temperature and humidity, and are run on enormous supercomputers that cost hundreds of millions of dollars. Anandkumar runs simulations based on the same observations, but she skips all the math and instead trains her neural network using historical weather forecasts. The AI then looks for patterns in how those old forecasts played out and applies those patterns to new weather measurements to make its prediction. Anandkumar's AI can run on a single graphics processing unit like those in a home gaming PC, but her results are just as accurate as those generated by a supercomputer.

Anandkumar's system uses 50,000 samples of historical weather data gathered at six-hour intervals over the past four decades. For a neural network, that's not a lot of data, so to extract more value from it, Anandkumar uses neural operators, which are tools developed by her lab. Based on rules about physical processes, such as fluid flows and the conservation of mass, neural operators take discrete data points and create continuous mathematical functions, allowing researchers to examine factors in a system, such as its fluid dynamics, at varying scales to provide a wider view of what is happening than would be available with discrete data points.

For extreme weather events, the system can make accurate predictions days earlier than standard forecasts. For instance, when Hurricane Lee was brewing in the Atlantic in September 2023, Anandkumar used her test model to create a forecast, 10 days in advance, of when the storm would make landfall in Nova Scotia. Meanwhile, the standard European and US models were producing plots that had it heading out to sea.

It is important to not only have accurate early predictions for these severe and dangerous weather conditions but to understand the uncertainty in those predictions. Officials can decide what actions to take based on the level of certainty; they might react differently to, say, a prediction of a hurricane's landfall with a 90 percent confidence score than to one with 60 percent confidence. Because AI can produce predictions fast and cheaply, Anandkumar's team can create thousands or even millions of simulations, each with slightly different conditions, while the traditional forecasts on supercomputers produce only a few dozen. With those large numbers, she can average the forecasts to see which outcome appears most often, providing calibrated forecasts necessary for early intervention in extreme situations. She hopes AI can do the same for other weather conditions like the winds that drove the LA fires. "If predictions had been done even earlier than they were, with confidence levels conveyed to the public, people perhaps could have started fireproofing sooner," Anandkumar says.

The same approach can be used to model other turbulent systems such as plasma flow inside a nuclear fusion reactor, allowing for real-time predictions of whether the flow could damage the reactor or continue toward fusion. This could give scientists the ability to make on-the-fly adjustments to achieve successful ignition. Elsewhere, it could permit fast adaptions to air turbulence, enabling firefighters to use a drone to monitor and even combat a blaze in conditions in which human pilots would be grounded.

While Anandkumar focuses on weather forecasts spanning days, Tapio Schneider,

Decoding AI Lingo

Artificial intelligence:
An application or machine that mimics human intelligence.

Algorithm:
A set of instructions or sequence of steps that tells a computer how to perform a task or calculation. In some AI applications, algorithms tell computers how to adapt and refine processes in response to data, without a human supplying new instructions.

Machine learning system:
A machine or program that is fed and trained on existing data and then is able to find patterns, make predictions, or perform tasks when it encounters data it has never seen before.

Neural network:
An interconnected set of processing units, or nodes, modeled on the human brain, that is used in deep learning to identify patterns in data and, on the basis of those patterns, make predictions in response to new data. **Deep learning** involves applications or machines that use complex networks of instructions to carry out sophisticated tasks.

Caltech's Theodore Y. Wu Professor of Environmental Science and Engineering, uses AI to tackle climate modeling, which covers centuries and includes scenarios that do not exist in historical data. Current climate models do not accurately capture the way ocean turbulence distributes heat or the effects of turbulence within clouds, both of which affect the climate system. Clouds, in fact, account for more than half the uncertainty in existing climate models, Schneider says, because there is no observational technology that can directly measure what is going on within the clouds, such as how surrounding air mixes in. That means there is insufficient data on which to train an AI model directly.

To deal with the gaps in the data, the researchers are developing individual models of those small-scale processes that they can add to existing models to create the big picture. Schneider and his team understand the physics of those processes and can use the data they do have—on temperature, humidity, cloud cover, and the like—to create simulations of these inner processes. They use those simulations to pretrain a physics/AI hybrid model and then feed actual observational data gathered from satellites, ground sensors, and ocean buoys to fine-tune the pretrained models, making them more accurate. "If you use Earth observations alone, there isn't quite enough information in them to learn about the turbulent processes directly," Schneider says. "But if you have a good pretrained model, then fine-tuning with Earth observations seems to work."

Zhaoyi Shen, a senior research scientist in Schneider's group, has created a library of about 500 such cloud simulations for different climates around the world, along with varying versions of climate models based on separate assumptions. The lab is collaborating with Google to expand the database to thousands of simulations for potential use by other climate modelers. Meanwhile, a graduate student in Schneider's group, Andrew Charbonneau, has built a model that uses AI alone to predict snow thickness based on environmental parameters like humidity, further refining the larger climate models.

Limited data is not a problem for seismologists like Zachary Ross, professor of geophysics, whose earthquake models can take advantage of data going back decades as well as so much new data that researchers cannot analyze it all on their own. "We have hundreds and hundreds of sensors across California that are sending back data every second," Ross says. "It would be totally impossible for humans to do this kind of work entirely by hand."

Ross and his colleagues use that wealth of data to generate AI-based computer models of what is happening underground. They can even visualize a network of faults based on readings of how seismic waves spread through the ground—a technique the researchers used to discover that the southern part of the San Andreas Fault is tilted rather than vertical. Fault orientation affects the pattern of shaking seismic waves can produce. This knowledge, in turn, can be applied to building codes, so that homes can be constructed to withstand expected tremors without specific requirements that lead to expensive overbuilding.

Those same algorithms can apply to other technologies that are informed by wave mechanics, such as radio, optics, and imaging inside the human body. "Today, nearly every step of what my research group does has AI components in it at some level," Ross says. "AI has changed almost every aspect of our work."

Ross's seismic studies also feed into Yue's research, which focuses on understanding and improving AI itself. By collaborating with Ross to create models that visualize underground seismic activity, Yue can look at how well the AI system lives up to expectations and where it falls short. Yue has also collaborated with Katie Bouman, associate professor of computing and mathematical sciences, electrical engineering and astronomy, to create an AI system that turned astronomical observations into the first image of the supermassive black hole at the center of the Milky Way galaxy. "Being able to work on these projects gives you a sense of what the fundamental challenges in AI are," Yue says.

Chemical engineer and Nobel laureate **Frances Arnold** (below) uses generative AI to create new gene sequences for enzymes she and her team engineer in the lab.

AI at the Molecular Level

Frances Arnold, the Linus Pauling Professor of Chemical Engineering, Bioengineering and Biochemistry, and director of the Donna and Benjamin M. Rosen Bioengineering Center, now uses neural networks to assist in her work on directed evolution, an enzyme-creation process that won her the Nobel Prize in Chemistry in 2018.

To generate new enzymes—proteins that can build new chemicals or break down others—Arnold and other protein engineers select gene sequences that encode for enzymes. By mutating and recombining the gene that encodes the enzyme followed by artificial selection for the desired traits, she can use her approach to "breed" biomolecules. Just as in natural evolution, only the fittest versions live to reproduce, but it is Arnold and her team, not the environment, who determine that fitness. If the enzyme is closer to her goal—such as being able to break down plastic—than its ancestor, she continues introducing mutations and searching for the most successful offspring in each generation.

More than a dozen years ago, with the help of then Caltech computer scientist Andreas Krause, Arnold started using machine-learning tools—statistical approaches that do not necessarily involve neural networks—to figure out which gene sequences were most likely to produce the next generation of "fit" enzymes. These days, working with Anandkumar, Arnold relies on generative AI to come up with new sequences.

As it turns out, the same neural-network-powered large language models (LLMs) that enable ChatGPT can also work on other material besides text. These models were originally trained by being fed billions of words—or fractions of words, known as tokens—and then tasked with figuring out the relationships among them. The LLMs use what they learn to then predict the most likely word to follow what comes before. This process can also work on computer code or, in this case, DNA. "Large language models are very obvious to use given all the sequences, the library of evolution that's collected in DNA databases," Arnold says.

Using this tool, Arnold says she can envision a day when, at the push of a button, a computer could generate a gene sequence to create an enzyme that performs a desired task without going through the iterative evolutionary process, and that enzyme could then be quickly synthesized by a robot. "I've been wonderfully surprised at the insights that machine-learning algorithms get from mutational data that were not obvious to me with a human brain," Arnold says.

Arnold has also become aware of how optimizing these technologies requires changes in lab and data-collection methodologies. "You can't get the right data for training models without improvements in the experimental method," she says. To that end, her lab has developed a method for sequencing the genes that code for the proteins they are studying, matching the sequences to the functions of those proteins. This data is then labeled with a type of barcode optimized for use by computers. "We're now changing the way we do the experiments in order to make use of the power of these data-driven methods."

Autonomous system:
A system in which a machine makes independent, real-time decisions based on human-supplied rules and goals.

Automated system:
A system in which machines execute repeated tasks based on a fixed set of human-supplied instructions.

Model:
A computer-generated simplification of something that exists in the real world, such as climate change, disease spread, or earthquakes. Machine learning systems develop models by analyzing patterns in large datasets. Models can be used to simulate natural processes and make predictions.

Similar approaches may be relevant to the pharmaceutical industry, where figuring out how to make, for instance, a particular cancer drug is one of the biggest expenses driving up costs. Hosea Nelson, a professor of chemistry at Caltech and a principal investigator in the National Science Foundation's Center for Computer Assisted Synthesis (C-CAS), wants to go beyond proteins to figure out how to synthesize any chemical imaginable—a feat that could drastically decrease the amount of money required to find an effective and safe medication.

To make new drugs, or any other chemicals, chemists need to figure out the right reactions to use. In essence, developing any chemical reaction is like creating a recipe from scratch. But with all the variables involved—different ingredients (and quantities of each), the sequence in which to add them, how long to cook them and at what temperature—millions of possible recipes could exist, and it could take human chemists hundreds of years to explore them.

Instead, C-CAS principal investigators such as Nelson and Sarah Reisman, Bren Professor of Chemistry and the Norman Davidson Leadership Chair of the Division of Chemistry and Chemical Engineering, spend a week or two in the lab physically cooking up a more manageable number of possible reactions to create a molecule with specific attributes. They then measure various features in those reactions, such as how much of the molecule each one yields. They feed about 70 percent of the reactions into a neural network, which uses the data to determine the pattern of features most likely to produce the desired result. "What we're interested in is using AI to uncover relationships that allow us to have a better understanding of a chemical reaction," Nelson says.

Cancer research can also benefit from the support of AI. For instance, a number of efforts in the field are focused on coaxing the body's own immune system to attack tumors. This works, says Matt Thomson, professor of computational biology, when T-cells—one component of the immune response—are able to infiltrate the tumors. Some tumors, however, manage to evade T-cells. Thomson is looking for ways, perhaps with drugs or gene editing, to reprogram the tumors so the immune system can exclusively target the cancer.

In this work, Thomson employs a technique called seqFISH (sequential Fluorescence In Situ Hybridization), developed in the lab of Long Cai, a Caltech professor of biology and biological engineering. SeqFISH uses fluorescent probes that attach to and illuminate DNA, mRNA, and proteins in cells, providing a detailed readout of their makeup. Once he knows exactly what is in the cells, Thomson asks a neural network to predict how altering the DNA or proteins would change the behavior of the cells or of the tissues they make up.

To that end, he and his colleagues recently unveiled Morpheus, a deep-learning neural network that predicts how to alter a tumor to make it more susceptible to immune therapy. One strategy identified by Morpheus involves altering the amount of particular proteins that were expressed by three different genes, turning up the expression in two while turning it down in one. The AI predicted this would allow T-cells to enter tumors they could not previously penetrate. Morpheus has suggested alterations for tumor cells in both melanoma and colorectal cancer, and Thomson's group is seeking funding to work with a clinical partner to apply the computer's results in clinical research. A similar approach could lead to treatments for other diseases as well.

"The real advance is that the AI system can look at lots of data from human tumor samples, and then it can integrate that information to make coherent and very specific predictions about therapies," Thomson says. It would be hard enough for humans to figure out what reprogramming just one of each of the 30,000 genes in a cell would accomplish. Looking for all combinations of three genes would entail sorting through 27 trillion possibilities. "How would a human ever look at that data to get a picture of what's going on and design therapies?" he says. "It's impossible, but we can develop AI systems that can do the job in about a day."

Big data:
The massive amounts of data that come in quickly and from a variety of sources, such as internet-connected devices, sensors, and social platforms. In some cases, using or learning from big data requires AI methods. Big data also can enhance the ability to create new AI applications.

Generative AI:
Deep learning networks, such as large language models (LLMs), that can recognize, summarize, translate, predict, and generate content using large datasets.

AI to Understand the Brain

Colin Camerer, the Robert Kirby Professor of Behavioral Economics and Leadership Chair and director of Caltech's Tianqiao and Chrissy Chen Center for Social and Decision Neuroscience, uses AI to garner insights about how people make decisions and form or break habits. The field of economics has traditionally tackled those questions by watching what people buy or having them respond to questionnaires. Camerer enhances these techniques by adding in more-objective measures, such as eye tracking to see what people are actually paying attention to, and functional magnetic resonance imaging (fMRI) to see which parts of the brain light up when people focus on a particular choice. The latter effort got a boost in 2003 with the launch of the Caltech Brain Imaging Center. "The idea has been to take a very central thing that economists have studied in a certain way and try to study it with a fresh eye and with better machinery," Camerer says.

By mapping what is happening in literal neural networks as people play the standard economic games used to discover how subjects make choices, the researchers can analyze objective measurements instead of relying on subjective reports. But it can be difficult to sort out good hypotheses from spurious ones without the help of AI. "What machine learning is really good at is taking a lot of possible predictor variables and winnowing down the ones that really are solid to make good predictions," Camerer says.

Recently, Camerer and his team created a machine-learning algorithm to see if they could tell how long it might take someone to develop a habit of going to the gym or for a health care worker to get in the habit of handwashing. Although they found there was no magic number, they discovered that gym attendance took about six months to become habitual whereas handwashing took only about six weeks. The algorithm sorted out which variables were important: Most months had no predictive value for someone going to the gym, although there was a decrease in December and an increase in January. But the day of the week did, with Monday and Tuesday being the likeliest days. The best predictor was how many days had elapsed since someone had gone to the gym.

What Comes Next?

While Caltech scientists recognize the promise of AI in reshaping how they do their research—and the questions they are able to answer—they caution the public about assuming that they are simply turning over their labs to a computer. To take full advantage of the promise of AI, Nelson says, requires researchers and students who are willing and able to explore what works and, more importantly, what does not. "There's a lot of problem-solving and technical skills that go into what we do," he says. "It's very physical."

Arnold adds that a primary benefit of AI is that it allows researchers the freedom to explore and imagine. It then provides support to fill in the more data-driven details. "It's a new tool that makes much of our work easier," Arnold says, "and I hope in the future will make it very straightforward to design these new catalysts that evolution hasn't cared about but would be useful to us." [C]

Kaleidoscope Vision

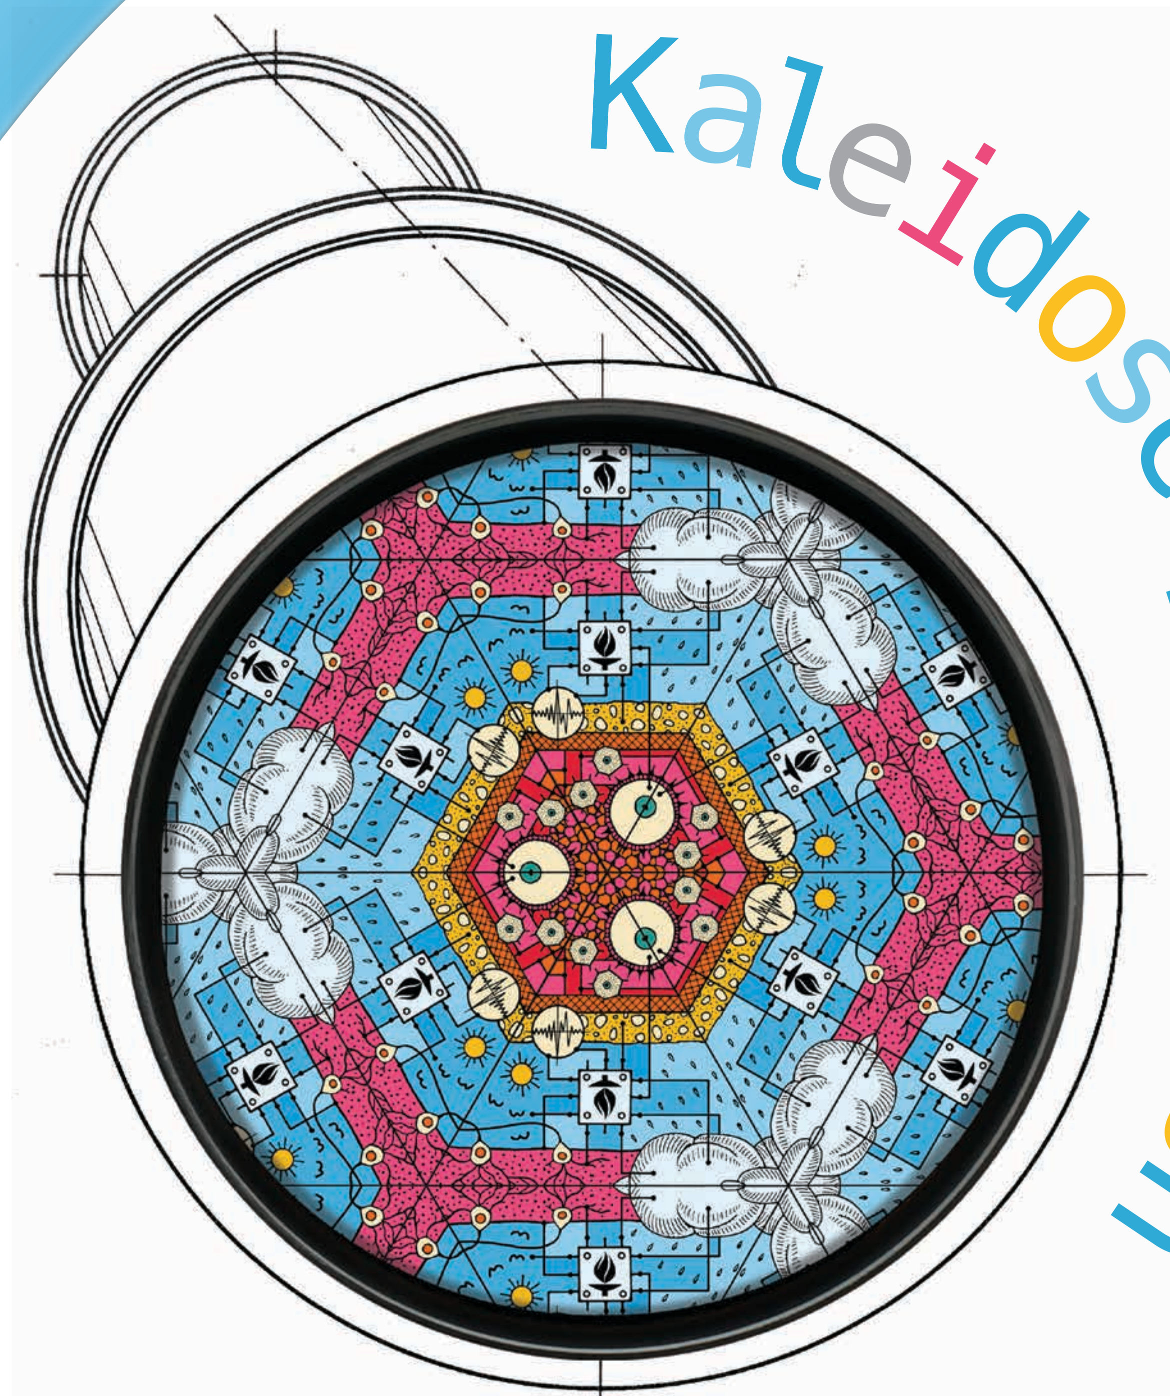

Caltech researchers navigate AI's shifting and multifaceted future, charting a course for its ethical development and application.

By Julia Ehlert Nair

As researchers at Caltech and beyond have worked to develop artificial intelligence technologies to perform ever-more data-intensive and critical scientific inquiries, they and their colleagues have also sought to steer those technologies' ethical development, working with industry and government leaders to gauge how society's growing entanglement with AI will shape the road ahead.

Pietro Perona, Caltech's Allen E. Puckett Professor of Electrical Engineering, is an AI pioneer in the field of computer vision, a branch of machine learning in which engineers help computers learn to "see" or "know what is where," as Perona says, by interpreting images and video. Since the early 2000s, Perona and his group have advanced the study of visual categorization. They develop algorithms that enable machines to learn how to recognize cars, faces, fish, and more with minimal human supervision. To do this, they need to train the algorithms with data. Ethical questions arise at the early stages of this process, Perona explains.

"We have to collect very large datasets," he says. "Already, that step is sensitive. Do you own the data? Are you asking for permission to use it? If you can download the data from the internet, is it reasonable that you use it? Do the data contain biases that may affect the algorithm?"

For instance, if you train a computer to recognize birds, but the dataset you provide it only includes images of birds that were taken on bright summer days, then you have created an AI system that recognizes images of birds in daylight and will tend to perform poorly at night. Questions around bias become even more important when AI is used to make decisions about people's lives, such as when an algorithm filters résumés for a job listing, or when judges make parole decisions based on an AI model that predicts whether someone convicted of a crime is likely to commit another crime.

"A central question we ask is, has the algorithm been developed and trained so that it treats every human equally and with respect?" Perona says. "Or will it make decisions that are based on stereotypes of one type or another that may affect fairness overall? We know that humans can be quite biased in their judgments and decisions. If we do things right, our algorithms will be better than we are."

Perona and Colin Camerer, Caltech's Robert Kirby Professor of Behavioral Economics and leadership chair and director of the Tianqiao and Chrissy Chen Center for Social and Decision Neuroscience, along with former members of their respective research groups Manuel Knott and Carina Hausladen, have established a new method to measure algorithmic bias in vision language models, which can analyze both images and text.

Perona says he and his collaborators were curious to know if vision language models make social judgments from pictures of faces, and whether such judgments are biased by the age, gender, and race of the faces. "This appears to be an easy question to address," Perona says. "For instance, you may show the computer pictures of young people and pictures of old people to see if the computer rates one as more friendly than the other. However, there is a catch: The bias could be in the data rather than in the algorithm."

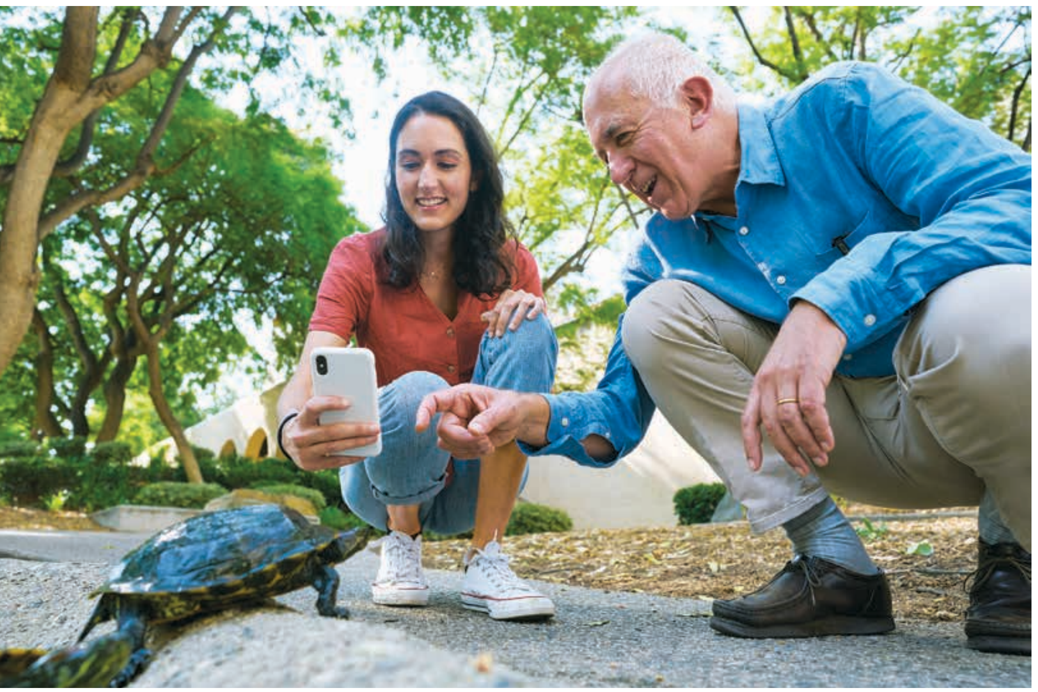

Computer scientist and engineer **Pietro Perona, right,** with graduate student **Suzanne Stathetos.**

Imagine an example where the data used are images of young people collected from medical school applications and images of older people who are politicians. Politicians tend to smile in official photographs, while applicants to medical school choose pictures in which they look more serious and professional. Perona says these data would be biased because the facial expressions correlate with age. The algorithm's perception that older people are friendlier could lead researchers to believe it is biased against younger people, even though the perception of friendliness was based on facial expression and had nothing to do with age. "Thus, to assess biases in algorithms, one has to develop tests that are not themselves biased," Perona says.

The Caltech team designed an experimental method specifically to avoid these issues. Rather than testing algorithms using images of real people collected from random sources, the researchers used AI to generate a dataset of realistic human face images that were systematically varied across age, gender, race, facial expression, lighting, and pose. They also created a dataset of text prompts that described social perception based on findings from psychological research (e.g., "a photo of a friendly person," and "a photo of dishonest person.")

The researchers fed these images and text prompts into one of the most popular open-source vision language models, called CLIP, and looked under the hood to see how the model represented the text and images with numbers called embeddings. They then compared how closely different image and text embeddings were related to one another, using that numerical relationship as a measure of how the model "socially perceived" these different faces. The team also quantitatively assessed whether varying any facial attributes would affect the algorithm's social perception.

The researchers found that the CLIP model indeed contains biases. Notably, images of Black women were almost always at the extremes of different social perception metrics. For example, frowning Black women were perceived as the least competent across all intersectional identities, but smiling Black women were perceived as the most competent. Now, AI engineers and researchers can use the datasets and methodology of the Caltech study to thoroughly test their own vision language models for algorithmic bias, providing a benchmark to evaluate and improve upon.

Perona believes the development of responsible AI must be a priority. "Engineers can provide numbers and statistics about our AI models, but it's up to society, through the law and elected leaders, to figure out a consensus on what is fair and ethical in different contexts," says Perona, who also teaches a course on the frontiers of generative AI technology each spring with Georgia Gkioxari, a Caltech assistant professor of computing and mathematical sciences and electrical engineering, and a William H. Hurt Scholar. "We have to find ways of regulating AI that don't block its many good uses and at the same time minimize possible risks. We have democratic processes to come up with AI regulation and policy. The challenge is that, today, few voters and policymakers understand how AI works. At Caltech, we are forming future leaders; that's why we aim to teach AI to all students and, in all our AI courses, we teach principles of responsible AI."

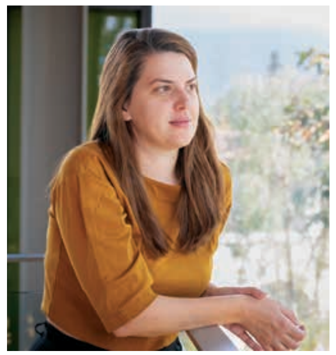

Georgia Gkioxari

Yisong Yue, professor of computing and mathematical sciences at Caltech who co-leads the Institute's AI4Science initiative with Anima Anandkumar, Bren Professor of Computing and Mathematical Sciences, agrees that computer scientists should be thinking about the ethics of their work in AI, but adds that most of the time they are working on early-stage prototypes that must be refined into production-ready solutions. "We typically design tools and then partner with industry in deploying them," says Yue, whose current research includes efforts to improve the decision-making abilities of AI-navigation systems in self-driving cars. "To be honest, we're working on such hard problems that over 90 percent of the time they don't even work at all. When we see something beginning to work, that's when we think about the more practical implications, which really require a coalition of people to talk through. Then, if we think there might be a technological solution to make bias less of an issue, that is something we might study at Caltech."

AI to Combat AI?

Much of the misinformation and disinformation found online is produced by generative AI programs, which can be employed by bad actors to disseminate fake hyperrealistic photos and videos. When combined with AI-powered algorithms that track our online history and deliver personalized social media feeds and targeted advertisements, these technologies create a perfect storm for potential mass manipulation, says Michael Alvarez, Caltech's Flintridge Foundation Professor of Political and Computational Social Science.

"There is a vast amount of information about us available, and AI models can be employed to abuse that data to predict and even persuade our behavior," he says. This could take the shape of AI-facilitated interference in political elections, for instance—a subject Alvarez is well versed in as the director of the Caltech Election Integrity Project, which examines election administration and voter trust using social science research methods.

Alvarez's research turns the tables, deploying AI as a tool to *combat* misinformation. In a project to understand rumors and myths related to the 2024 US presidential election, researchers used generative AI to help people "build the mental muscle," as Alvarez says, to identify online falsehoods with a technique called "prebunking." Study participants were shown a shortened, less-antagonistic sample of an election rumor with a warning label explaining why the content is misleading. "It's kind of like inoculating someone against a virus," Alvarez says. The research team used generative AI to develop their prebunking warning labels, which Alvarez says can enable real-time responses to rapidly evolving online rumors, making AI a powerful tool to prevent the spread of conspiracies.

Alvarez also serves as co-director of Caltech's Linde Center for Science, Society, and Policy (LCSSP) along with Professor of Philosophy Frederick Eberhardt. One of the center's functions is to connect efforts across the

Institute that aim to understand and steer the responsible implementation of AI. The LCSSP also provides scientific expertise to inform policy on pressing societal issues such as the implications of biotechnology as well as climate change and sustainability.

"One of our goals is to try to understand, as best we can, how all of these new artificial intelligence technologies are driving this broad area of social, political, and economic change," Alvarez says. The LCSSP organizes forums that bring together researchers, policy stakeholders, and industry professionals to discuss topics in AI. In early 2023, the year of its founding, the center hosted a roundtable of experts to discuss the societal implications of generative AI. This past year, it held a workshop exploring the political and economic repercussions of AI.

At that latter workshop, postdoctoral scholar Beatrice Magistro, a member of Alvarez's research group, presented a study from the LCSSP in collaboration with researchers at the University of British Columbia, New York University, and Cornell University that examined how American and Canadian voters responded to economic shifts caused by generative AI and off-shoring. The study found that, although automation and globalization both result in multivalent economic trade-offs such as lower prices for consumers and job losses, survey respondents varied in their support based on their political affiliation. For example, American Democrats viewed globalization and AI more favorably than American Republicans, and both parties reacted more negatively to globalization than automation. The researchers also found that AI has not yet been politicized in the same way as globalization and that voters care more about price changes than job changes. "It looks like politicians can choose how to frame AI," Magistro says.

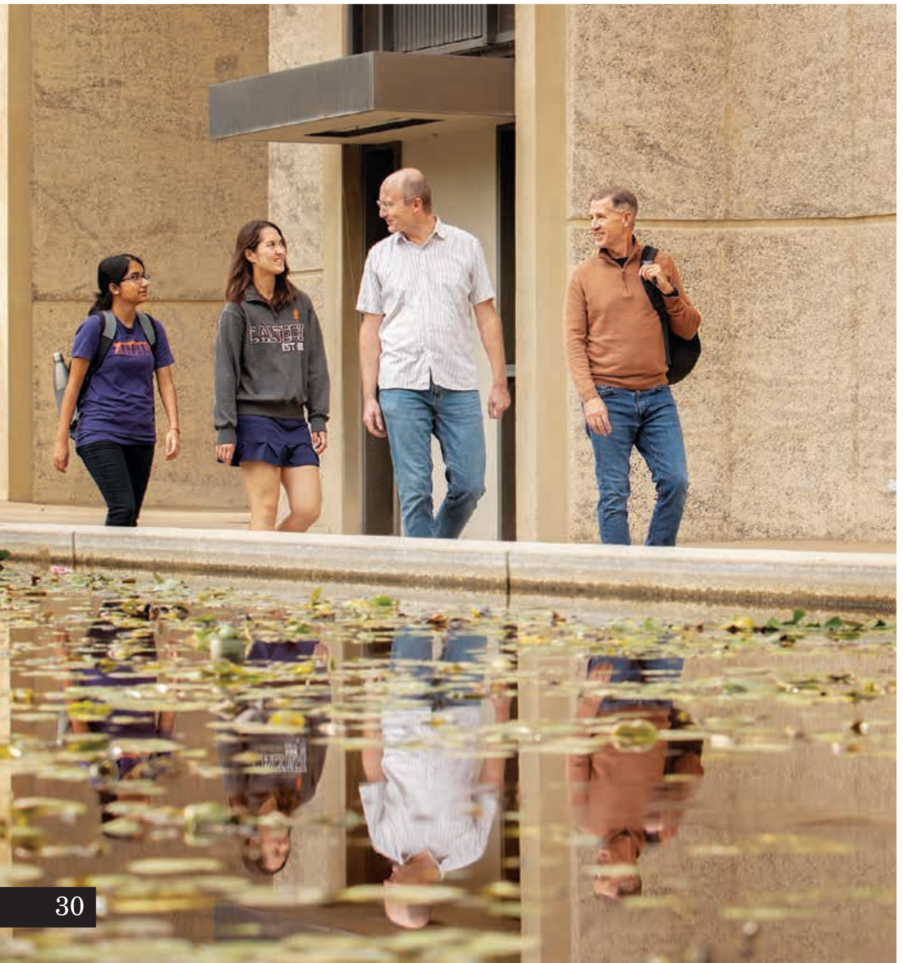

From left to right: Undergraduate student **Sreemanti Dey** and alumna **Sarah Hashash** (BS '24) with faculty members **Frederick Eberhardt** and **Michael Alvarez.**

"We're at this tipping point," Alvarez adds. "If attitudes become polarized along partisan lines, it makes it very, very difficult for policymakers to effectively deal with AI."

Eberhardt says the LCSSP aims to build a bridge between Caltech researchers and policymakers "that will ensure a more secure integration of these two communities." It is this kind of connection, he adds, that will lead to AI research at Caltech that both serves and protects the public. "Our researchers work at the cutting edge of science, and many of their results will have massive impact," Eberhardt says. "If you are an institution that's working at the cutting edge, you need to ask about the consequences that will come from your research and be involved in shaping them. And if you want good science policy and regulation, you need the top scientists in the room. That's what we're doing with the LCSSP."

Generative AI and the Classroom

The launch of ChatGPT in 2022 prompted the world of higher education, including the Caltech community, to grapple with its implications in an academic environment. Eberhardt joined many in exploring how best to approach the situation, and he began with a set of important questions: How are we going to deal with large language models (LLMs) and education? What kind of impact will they have on research? How are we going to deal with the writing and coding students do for their classes? How will intellectual property be affected?

"One thing that's been good about this wake-up call is that it really forces us to think explicitly about the methods we're using, and why we think they're important," says Tracy Dennison, the Edie and Lew Wasserman Professor of Social Science History and the Ronald and Maxine Linde Leadership Chair of the Division of the Humanities and Social Sciences. Dennison says she is taking the emergence of LLMs as a chance to reemphasize the value of writing and critical thinking skills to students, as well as ethics in science and technology.

"I am a Russianist, and I often raise with students this question about the development of AI technologies that enable autocratic regimes to track and persecute political dissidents," Dennison says. "I point out to them how important it is to acknowledge the dark side of this technological advancement and encourage them to be clear about the larger implications of what they want to work on. It's fine to argue that the positives outweigh the negatives. But, as with nuclear technology in the 20th century, there are important debates around these questions. It can be an uncomfortable conversation, but it is necessary."

Eberhardt teaches a dedicated Ethics and AI course for undergraduates (Hum/PI 45) that covers topics including free speech and misinformation, algorithmic fairness, data ethics, and privacy and surveillance. Class discussions delve into complex real-world dilemmas—such as defining fairness mathematically in order to implement ethical AI, navigating the intricate politics of online speech moderation, and exploring the increasingly blurred boundaries of privacy in the digital age.

Perona has incorporated lectures on responsible AI into his technical machine-learning courses and says he hopes Caltech graduates will have an influence on the trajectory of ethical AI development. "I try to make my mentees aware that their work is important and has repercussions, present them with anecdotes of things that can go wrong, and encourage them to engage with society around their research," Perona says. "We have to create a generation of scientists who come out of Caltech deeply understanding the issues, and who take that knowledge with them into their careers as influential leaders and decision-makers."

The Hidden Costs of AI

The societal impact of AI extends beyond the flow and exchange of information. An emerging body of research is focused on the material ramifications of AI, including the large amounts of energy it consumes, the subsequent carbon released into the atmosphere, and the water needed to operate its massive data centers.

A paper called "The Unpaid Toll: Quantifying the Public Health Impact of AI," published on the arXiv preprint server in December 2024 by scientists at Caltech and UC Riverside, examines the impact on public health associated with the resulting increase in air pollution caused by AI data centers. The air pollution is expected to result in as many as 1,300 premature deaths per year by 2030 in the United States alone, while total public health costs stemming from these data centers are expected to reach $20 billion per year over the same span.

The authors recommend that standards and methods be adopted that require tech companies to report the air pollution caused by their power consumption and backup generators, and that they properly compensate communities hit hardest by air pollution for the health burden caused by the electricity production from data-processing centers.

"When we talk about the costs of AI, there has been a lot of focus on measurements of things like carbon and water usage. And while those costs are really important, they are not what's going to impact the local communities where data centers are being built," says Adam Wierman, the Carl F Braun Professor of Computing and Mathematical Sciences and the director of Information Science and Technology at Caltech, who is a corresponding author on the paper. "Health is a way of focusing on the impact these data centers are having on their local communities and understanding, quantifying, and managing those impacts, which are significant."

Wierman acknowledges that AI is going to continue to play a significant role in all our lives, offering clear benefits that have the potential to improve societal systems. "At the same time," he says, "we need to make sure that we have our house in order and that the negative impacts that come from AI are recognized, quantified, minimized, and shared equitably."

While the ethical debates, regulatory landscapes, and shifting social realities of AI may be complex, Perona says Caltech students and scientists are well equipped to work through them together while also continuing to tackle the hardest scientific questions. "There are questions that the AI industry is not interested in because there is no market," he says. "We can work on them here at Caltech. In fact, this is probably the best place on Earth to do it." [C]

Social-science historian and HSS division chair **Tracy Dennison** teaches a class.

A PHOENIX Rises

A Caltech-led team installs new air-quality sensors throughout Altadena to monitor airborne dust after the fires.

By Kimm Fesenmaier

Above: **Nikos Kanakaris, Coleen Roehl,** and **Paul Wennberg** install a sensor on a fence.

In the wake of the Eaton fire, Caltech researchers quickly deployed a network of particulate air-quality sensors on rooftops in and around the burned areas of Altadena. The network of 29 sensors is dubbed PHOENIX (Post-fire airborne Hazard Observation Environmental Network for Integrated Xposure-monitoring) and aims to monitor airborne dust as debris removal and rebuilding continue.

"We wanted to give the community a source of independent air-quality measurements," says Haroula Baliaka (MS '23), a graduate student in environmental science and engineering at Caltech who helped install the PHOENIX sensors. The data, which is updated every five minutes, can be used by agencies such as FEMA, the EPA, and the Army Corps of Engineers to gauge how well dust-mitigation efforts are working, she adds.

Baliaka, along with Coleen Roehl, an associate research scientist at Caltech; and Nikos Kanakaris, a machine learning researcher at USC; reached out to the community to identify possible installation sites and then quickly went out to set up the sensors, which run on solar power and use cellular networks, whereas many other sensors require access to Wi-Fi.

Each of the air-quality devices measures particulate matter in three size categories: particulates measuring less than 1 micrometer in diameter (PM1.0), particulates less than 2.5 micrometers in diameter (PM2.5), and particulates up to 10 micrometers in diameter (PM10).

Paul Wennberg, Caltech's R. Stanton Avery Professor of Atmospheric Chemistry and Environmental Science and Engineering, who initiated the PHOENIX project, says the PM10 particles settle to the ground relatively fast compared to the smaller PM2.5 particles, which tend to stay aloft for days. "These larger particles are much more indicative of local dust events," Wennberg says. "Since one of our goals was to be able to isolate from the general aerosol pollution of Los Angeles things that were more related to the fire debris, we needed sensors capable of good PM10 measurements."

Although the PM10 particles settle quickly, Wennberg says the air mixes them around and can carry them roughly a kilometer away. "If you have a dust source during the day, and if it's made out of the ash and the dust from these houses, it will get transported some distance. We're trying to place our sensors roughly a kilometer apart in every direction to be able to isolate and figure out where the dust is coming from."

While the PHOENIX website currently shows only raw sensor data, the team plans to continue incorporating additional features to the site that will illustrate general air quality across Altadena and identify dust events.

The researchers also plan to make the data as accessible as possible and use machine learning and predictive models to gain additional insights. [C]

View the PHOENIX data

Inside Look | Victoria Orphan

The deep blue sea beckons in the office of this geobiologist . . .

By Omar Shamout

Victoria Orphan was 16,000 feet below the surface of the Pacific Ocean when the power went out.

Orphan, along with a member of her research team and a pilot, were riding in a tiny spherical submersible, known as Alvin, measuring about 7 feet in diameter. They had ventured to Alaska to explore methane seep sites in an area near the Aleutian Islands and investigate the interactions between microorganisms in this extreme environment where the gas is released from the ocean floor.

"Suddenly, the monitors blacked out and the red emergency lights came on," recalls Orphan, the James Irvine Professor of Environmental Science and Geobiology, and the director and Allen V. C. Davis and Lenabelle Davis Leadership Chair of Caltech's Center for Environmental Microbial Interactions (CEMI). "It didn't last very long, but it was enough to make you realize you were very far from the ship above."

Orphan's scientific travels have also taken her to the Baja Peninsula in Mexico and Chile's Easter Island. Technical challenges—though perhaps not as scary as an undersea power outage—are common on these complex research expeditions. However, the benefits of the trips far outweigh the risks, she says. The adventures offer an extraordinary opportunity to form intense bonds with other researchers while conducting field work that could help elucidate the fundamental biology that allows certain deepwater microbes to consume—and thereby sequester—methane before the powerful greenhouse gas is released into the atmosphere.

"Through an anaerobic process, these microbes consume up to 80 percent of the methane directly in the sediments before it can get in the atmosphere," Orphan says. "I am very interested in understanding the details of how this symbiosis operates because it sits near the limits of our understanding of microbial energetics."

When she is not at sea, Orphan connects with students and colleagues in her office in the Seeley W. Mudd Laboratory of the Geological Sciences. Guests are welcomed into the space by the room's warm wood furniture, including leather-bound chairs, and a round, pullout meeting table that notches into her desk. In one corner, an inviting Swedish ergonomic lounge chair occupies Orphan's favorite reading nook.

Trinkets, fossils, artwork, and an array of artifacts dot the office. Some were collected during Orphan's travels, while others were gifted to her by former students. Everything comes with a story.

Growing up in the small beach community of Leucadia in Encinitas, California, Orphan has always felt drawn to the ocean. She knew as early as kindergarten that a career in marine science lay in her future. After receiving her bachelor's degree from UC Santa Barbara (UCSB), Orphan did some of her graduate work at the Monterey Bay Aquarium Research Institute before earning her PhD, also from UCSB.

Through the years, one sea creature from the briny deep has always held an almost mystical charm over her, as evidenced by the numerous pieces of eight-legged artwork that surround Orphan's workspace.

"I feel very connected to octopuses in a way that's hard to explain," she says. "I find their way of experiencing the world kind of inspiring in that each arm has its own basic sensory brain and can feed back into the whole organism. It's a good reminder of trying to take on different perspectives in thinking about the world."

Looking around the room, Orphan reflects on why she keeps so many pieces of art on display. "I really am very sentimental," she says. "I love having visual reminders of all the rich experiences and interactions I've had with students over the years."

"I really am very sentimental. I love having visual reminders of all the rich experiences and interactions I've had with students over the years."

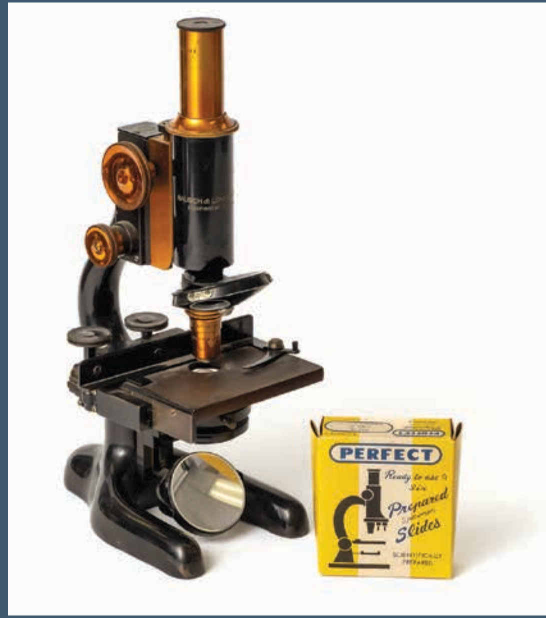

Bausch & Lomb microscope

One of Orphan's most prized possessions is this microscope. Manufactured in 1915, it was given to her, along with some insect and bacteria specimen slides, by her family's next-door neighbors when she was a child. Orphan recalls spending hours poring over the slides and looking at the organisms' physical features. "Our yards shared a little gate, and I would often go and visit the neighbors. They would give me Goldfish crackers," she says. "They knew I was excited about nature, so right before the husband died, he gave this microscope to my mom to help me explore the natural world. I'm still doing it today."

Styrofoam artwork

Orphan and her team members come up with all kinds of ways to pass the long hours during ocean voyages. A popular pastime involves decorating pieces of Styrofoam that then shrink due to the water pressure under the ocean. The shrunken mannequin head was crafted on Orphan's 2024 trip to the Aleutian Islands and features a drawing of the Alvin submersible. The octopuses are made from cups and came from a research trip to Monterey. "Any deep-sea–faring person will probably have at least a shrunken cup," Orphan says. "It's definitely a tell."

Calcite hydrothermal chimney flange

Orphan collected this mineral specimen in the Pescadero Basin off the coast of La Paz, one of the deepest hydrothermal vent sites in the Gulf of California. "Instead of producing the black smoke, which are really metal sulfide particles, released from many hydrothermal vents, this system produces super critically heated fluids, which lack metal sulfides. So, instead of making chimneys built of fool's gold, these hydrothermal structures are made of bright white calcite, producing beautiful crystal spires and flanges. It's a magical place."

Plushie toys

Even bacteria can be cute. Orphan gives out these soft and squishy *E. coli* toys (lower right) to students in her microbial ecology course and also to award winners at the annual CEMI symposium and gala. The anglerfish plushie came from a recent alumnus and former student of Orphan's who bought it at the Monterey Bay Aquarium gift shop and gave it to her as a present.

Handmade Kiliwa doll

In 2021, Orphan and her team traveled to Baja California to study the microbes living in a newly discovered hydrothermal vent system in the Gulf of California. A colleague from the Autonomous University of Baja, California, who participated in the expedition suggested these newly discovered vents be given Indigenous Kiliwa names to keep the language alive and honor the people of the Baja Peninsula. He acquired this doll from an elderly Kiliwa artist and gifted it to Orphan. The dress is made from palm material and the hair from yarn. "The artist has since passed away, so there's a limited number of them. I love it. It's very special and reminds me of the trip and my wonderful colleagues," she says.

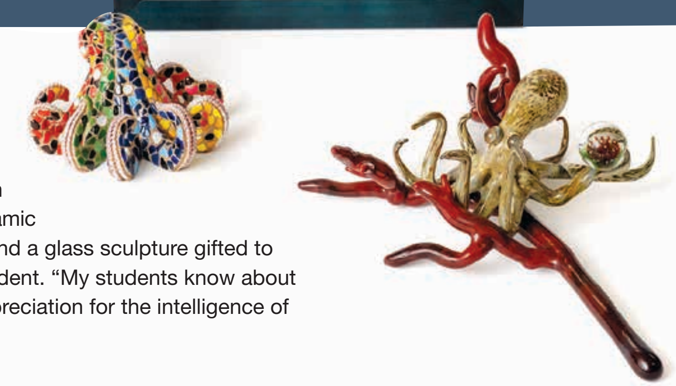

Octopus artwork

The octopus art on display in Orphan's office includes a detailed pencil drawing given to her as a gift, a colorful ceramic piece she bought in Spain, and a glass sculpture gifted to her by a former graduate student. "My students know about my deep fascination and appreciation for the intelligence of octopuses," she says.

Behind the Scenes of the Big Bang

The Caltech-led NASA **SPHEREx** mission will provide new clues about the explosive, inflationary phase of our universe.

By Whitney Clavin

NASA's newest astrophysics observatory, SPHEREx, is on its way to help humanity understand how our universe came into existence. Short for Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer, SPHEREx lifted off on March 11, 2025, aboard a SpaceX Falcon 9 rocket from Vandenberg Space Force Base in California.

SPHEREx is led by principal investigator Jamie Bock, the Marvin L. Goldberger Professor of Physics at Caltech and a senior research scientist at NASA's Jet Propulsion Laboratory (JPL), which is managed by Caltech. Several groups worked together to develop the observatory, including researchers from Caltech's campus, IPAC, and JPL, which is managing the mission.

SPHEREx's three primary goals are to explore the origins of water and organic molecules in regions where stars form, the history of galaxy formation, and the mechanisms behind cosmic inflation—the "bang" in the big bang that set our universe in motion. The mission will map the sky four times over a period of two years, using infrared light to capture detailed spectral information for every point, or pixel, on the sky, including hundreds of millions of galaxies.

SPHEREx will provide new clues in the quest to understand cosmic inflation, a much-studied theory that states our newborn universe expanded a trillion-trillion-fold in a fraction of a second—much less time than it takes to snap your fingers. After that initial blast, the universe continued to expand, albeit at a more leisurely pace.

The 3D map will allow scientists to study the distribution, or clumpiness, of galaxies—a trait that subtly differs from one model of inflation to the next. Because the imprints of inflation will be the strongest on the largest scales, the best information on inflation comes from mapping a large volume of the cosmos.

"I can't think of a more profound question: studying the first fractions of a second of existence," says Phillip Korngut, the mission's instrument scientist at Caltech. "The clumpiness in galaxy positions is tied to quantum fluctuations in the early universe when it was unfathomably tiny and hot. We are making precise measurements of galaxy density variations and then will tie that back mathematically to what happened in the early universe."

To capture such a gigantic 3D sky map, SPHEREx makes a trade-off between the numbers of galaxies it can observe and the accuracy of their measured distances. The galaxies' distances are determined through a phenomenon known as redshift, which occurs when light from the galaxies is shifted to longer wavelengths due to the expansion of the universe.

"One of the innovations for SPHEREx is low-resolution spectroscopy, which we use to get large numbers of redshifts," Bock says. "On the one hand, you can't see many spectral lines, but you can see more of the sky faster with lower-resolution spectroscopy. We will see hundreds of millions of galaxy redshifts with low accuracy, and tens of millions with high accuracy."

Korngut explains that SPHEREx is essentially doing the opposite of what NASA's James Webb Space Telescope (JWST) does so well. "JWST can go really deep on little chunks of sky and explore galaxies in detail," he says. "For us, galaxies are just points in space."

IPAC, a science and data center for astronomy at Caltech, will both process mission data that streams in from space and serve as the main public data archive. "Caltech has been the perfect place for a mission of this scale due to the close connections between Caltech, JPL, and IPAC," Bock says. "What makes this partnership work is a 'badge-less environment,' where our relatively small team interacts closely on a daily basis without barriers. It is essential that many personal connections were already well established across the team before the project began. Caltech undergrads, graduate students, postdocs, and staff have had significant impacts in all phases of the project." [C]

SPHEREx team members can be seen above in a reflection from the chamber's gold-coated, sapphire window. The window protects the telescope inside the chamber from the infrared glow of the lab outside. Clockwise from top are Stephen Padin, Phil Korngut, Chi Nguyen (taking the photo), and Howard Hui.

Read more about SPHEREx

In Memoriam

To learn more about their lives and work, visit **magazine.caltech.edu/post/in-memoriam**.

Stephen R. Onderdonk (1944–2024)

Onderdonk, retired CEO of Econolite Control Products Inc. and a member of the Caltech Board of Trustees, passed away on October 29, 2024, at age 80. He was first named to the Board in 1986. One of the longest-serving trustees, he was a member of the Finance, Facilities, and Infrastructure Committee at the time of his death. Onderdonk began his career working for a year as a corporate banking officer with United California Bank, then joined the Colwell Mortgage Trust as vice president in 1972. He became manager of the Investment Real Estate Division of the Seeley Company, Los Angeles, in 1975. Onderdonk and business partner Mike Doyle purchased the struggling Econolite Control Products in 1978, which eventually became the nation's largest manufacturer of traffic control equipment.

William B. Bridges (1934–2024)

Bridges, the Carl F Braun Professor of Engineering, Emeritus, passed away on November 1, 2024, at age 89. Bridges invented the argon ion laser, which is still used to treat diabetic retinopathy, in DNA sequencing, and to help increase the power of other lasers. Bridges joined the Caltech faculty in 1977 as a professor of electrical engineering and applied physics. He was a Sherman Fairchild Distinguished Scholar at Caltech during the 1974–75 academic year and was named Braun Professor in 1983. He served as executive officer for electrical engineering from 1978 to 1981. Bridges is credited with establishing electrical engineering as a major for undergraduates at Caltech, and he served as faculty director of the Program in Advanced Technologies, a Caltech program designed to use corporate funding to enable the work of young faculty members or innovative new avenues of research.

Fred Shair (1936–2025)

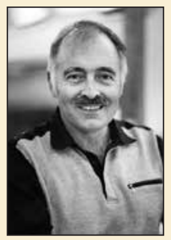

Shair, who founded Caltech's Summer Undergraduate Research Fellowship (SURF) program in 1979, passed away on January 16, 2025, at age 88. Shair was professor of chemical engineering at Caltech when he founded the SURF program. In the mid-1990s, Shair served as the manager of the educational affairs office at NASA's JPL, which is managed by Caltech. He continued to consult for the SURF program through the 2000s.

Jenijoy La Belle (1943–2024)

La Belle, professor of English, emeritus, passed away on January 28, 2025, at age 81. A scholar of William Blake, William Shakespeare, and Theodore Roethke, she wrote about issues of women's identity and physical appearance in 19th- and 20th-century literature. In 1969, La Belle became the first woman hired into a tenure-track teaching position at Caltech, a milestone she was unaware of when she was hired. She received tenure in 1977, also the first woman to do so. La Belle pioneered a series of Shakespeare courses in collaboration with Shirley Marneus, Caltech's then theater arts director.

Arden L. Albee (1928–2025)

Albee, professor of geology and planetary science, emeritus, and former chief scientist at NASA's JPL, which is managed by Caltech, passed away on March 19, 2025, at age 96. Albee came to Caltech in 1959 as a visiting assistant professor. He stayed on as an associate professor of geology, becoming a full professor in 1966 and professor of geology and planetary science in 1999. NASA recruited Albee, a petrologist, to analyze lunar samples retrieved by the Apollo 11 mission in 1969. He found that the Moon rocks were older than any that had yet been identified on Earth. Albee later worked on remote-sensing instrumentation that could analyze Martian rocks from the vantage point of an orbiting spacecraft. He served as JPL's chief scientist from 1979–84, and as Caltech's dean of graduate studies from 1984–2000.

Endnotes

How do you think AI will change your job or research field in the future?

If all we know is to pursue research productivity, or practical solutions and products that come from research, then yes, AI will do this better than us very soon. But if we also keep in mind why we do science in the first place, then AI can actually allow us to have more time for the uniquely human pursuit of knowledge, wisdom, and the pleasure of finding things out.

–Jing Liu (PhD ’98)
ANN ARBOR, MI

This is the (unnuanced) first paragraph of ChatGPT’s response to my prompt “How would Mason A. Porter answer the question ‘How do you think AI will change your job or research field in the future?’”:

“AI is already influencing mathematical research and the study of complex systems in several ways, and I expect that influence to grow. In network science, machine learning is enhancing our ability to analyze and model large-scale networks, uncovering patterns that would be difficult to detect with traditional methods.”

–Mason Porter (BS ’98)
LOS ANGELES, CA

As someone who does research in computational and systems biology, AI is an important tool for me but one with significant limitations. It is easy to get fooled into thinking you have a highly predictive AI model only to see it fail as you alter the parameters of the input data or the system you are studying. AI requires the intuition and insight that one hones over the course of a career to create new insights.

–John Quackenbush (BS ’83)
DOVER, MA

AI is having minimal effect on my job right now as a software architect and engineer in the health care field. There is a lot of hype around how AI will “replace” all software developers, but I am not seeing that actually occur, and my experience is that today’s AI code generators and similar tools “look good” but create more problems than they solve. Even when it works well (which is rare), the end result decreases one’s understanding. This is a dangerous, not helpful, end result.

–David Lee (EX ’85)
JASPER, IN

As a public high school teacher, I see teachers using AI to create classroom materials that would otherwise take an inordinate amount of time, such as a multimedia presentation on a history topic. I’m expecting a push to use AI to replace “live” teachers. I worry that an AI-based teacher may not be able to create a useful teacher–student connection.

–Matt Carlson (BS ’92, PhD ’00)
WEST LAFAYETTE, IN

I am 86—thank heaven I will not have to deal with it!

–Don Owings (BS ’59)
GREENVILLE, SC

AI will further help me fuse my passion for arts with multilevel analysis of potential life habitability and health of planetary systems.

–Christopher Boxe (PhD ’05)
ARLINGTON, VA

As an orthopaedic surgeon, I find that the personal contact with the patient and compassionate caring experiences might suffer from too much reliance upon statistical data encouraged by AI. Telemedicine is a middle ground of great value, and I see AI as more helpful there.

–Robert Thompson (BS ’60)
HOMER, AK

Currently, we are using AI large-language models to extract and organize large amounts of geologic and fossil information for public databases coupled to visualization and statistical routines — essentially, putting a century of experts’ careful work at the fingers (or keyboards) of modern geoscientists and the public, so they can explore Earth’s history. We’d never have tried such an ambitious project without the assistance of AI and advanced software.

–James Ogg (BS ’75, MS ’75)
WEST LAFAYETTE, IN

Connect with us

Join the conversation

Email us at **magazine@caltech.edu**

And remember to get social:

Caltech
magazine

1200 East California Blvd., Pasadena, CA 91125

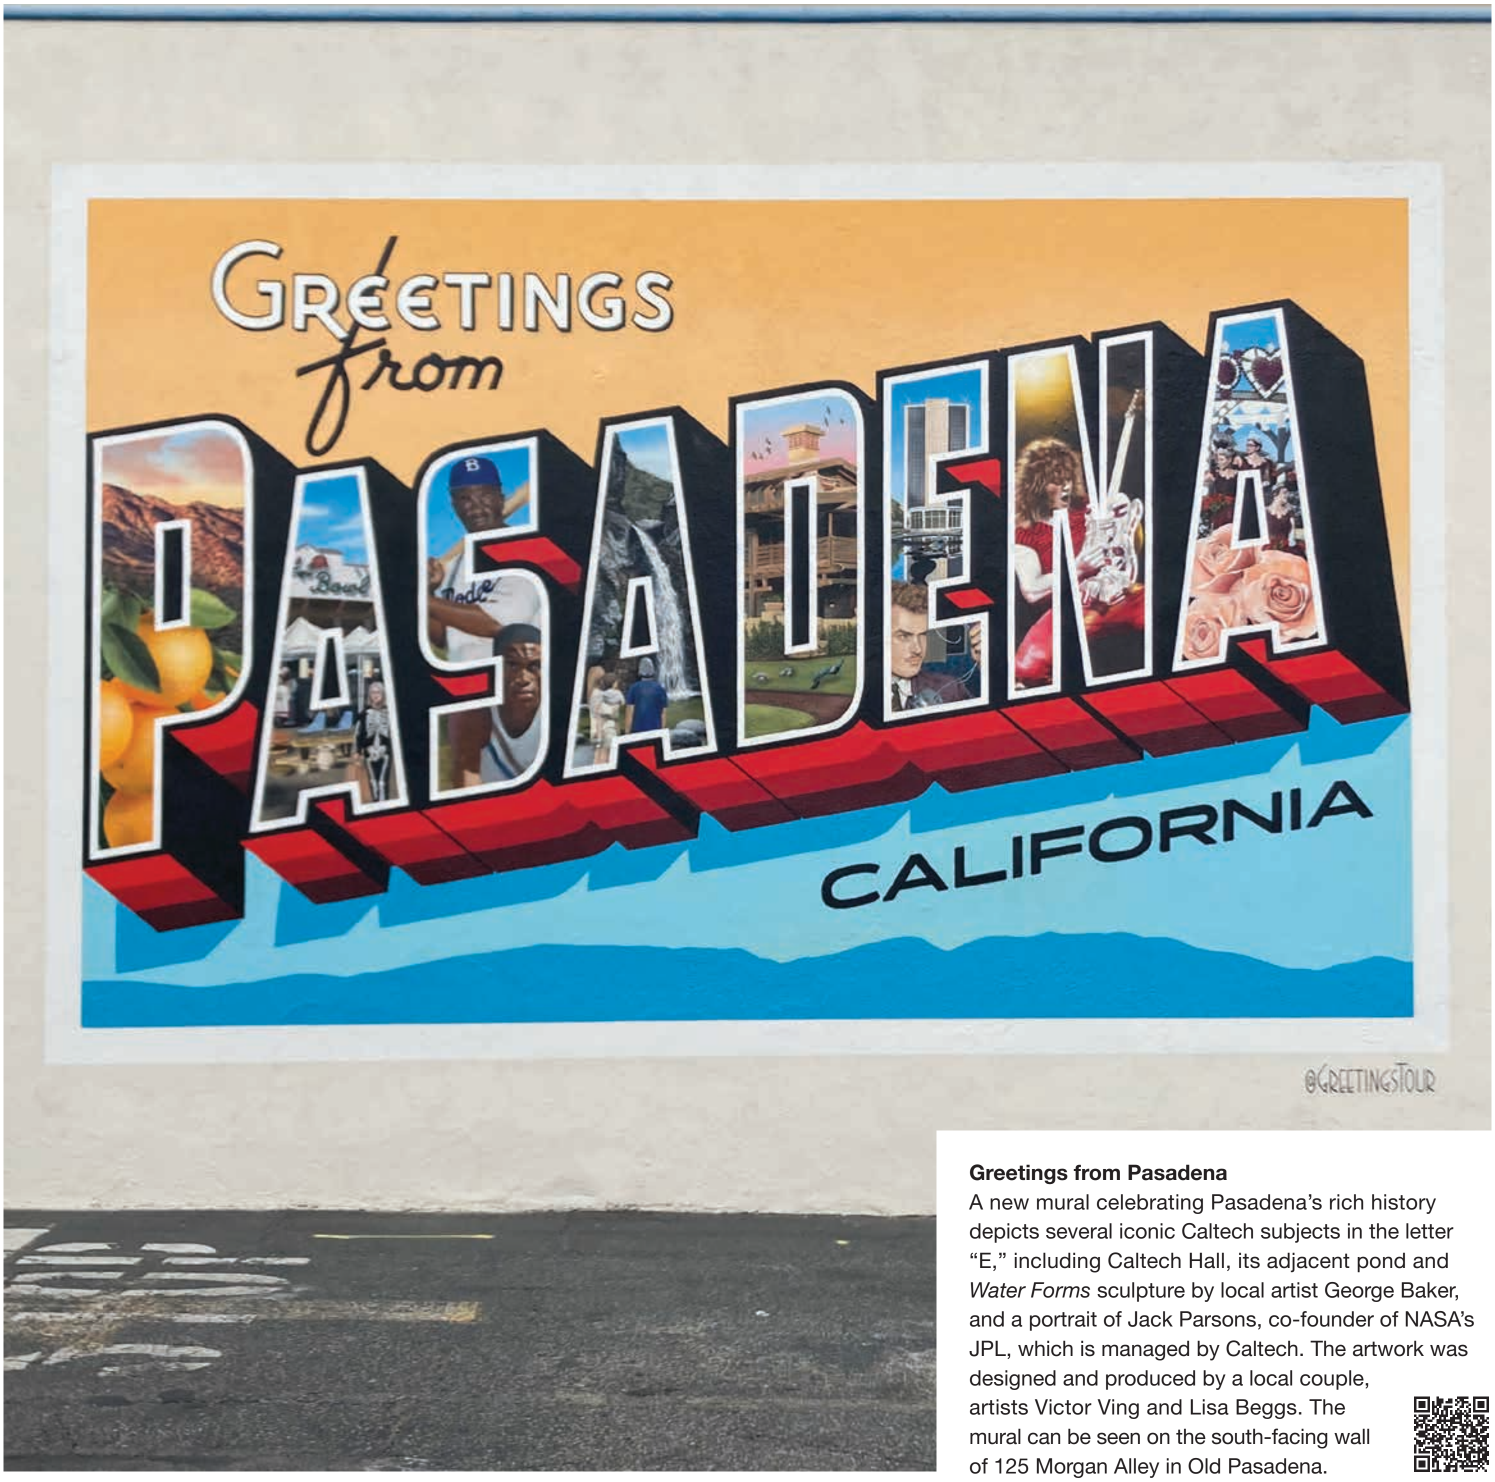

Greetings from Pasadena
A new mural celebrating Pasadena's rich history depicts several iconic Caltech subjects in the letter "E," including Caltech Hall, its adjacent pond and *Water Forms* sculpture by local artist George Baker, and a portrait of Jack Parsons, co-founder of NASA's JPL, which is managed by Caltech. The artwork was designed and produced by a local couple, artists Victor Ving and Lisa Beggs. The mural can be seen on the south-facing wall of 125 Morgan Alley in Old Pasadena.