S T A N F O R D
M E D I C I N E
Issue 3 / 2023

special report

AI EXPLODES
Taking the pulse of artificial intelligence in medicine

Cautiously optimistic
Leaders of Stanford's responsible health AI initiative look ahead

Targeting the tumor's driver
Gene testing and AI to halt breast cancer

Who's training whom?
A physician's surprising encounter with ChatGPT

A way to fill AI's massive energy demand
Chips that work like neurons

Will the brain always be best?
Four neuroscientists consider the question

plus

Out of darkness, hope
A family devastated by cancer plants seeds to vanquish the disease

Fast-forwarding drug development
Clinical trials in a dish

When food is more than just food
A neurosurgery trainee contemplates her father's final days

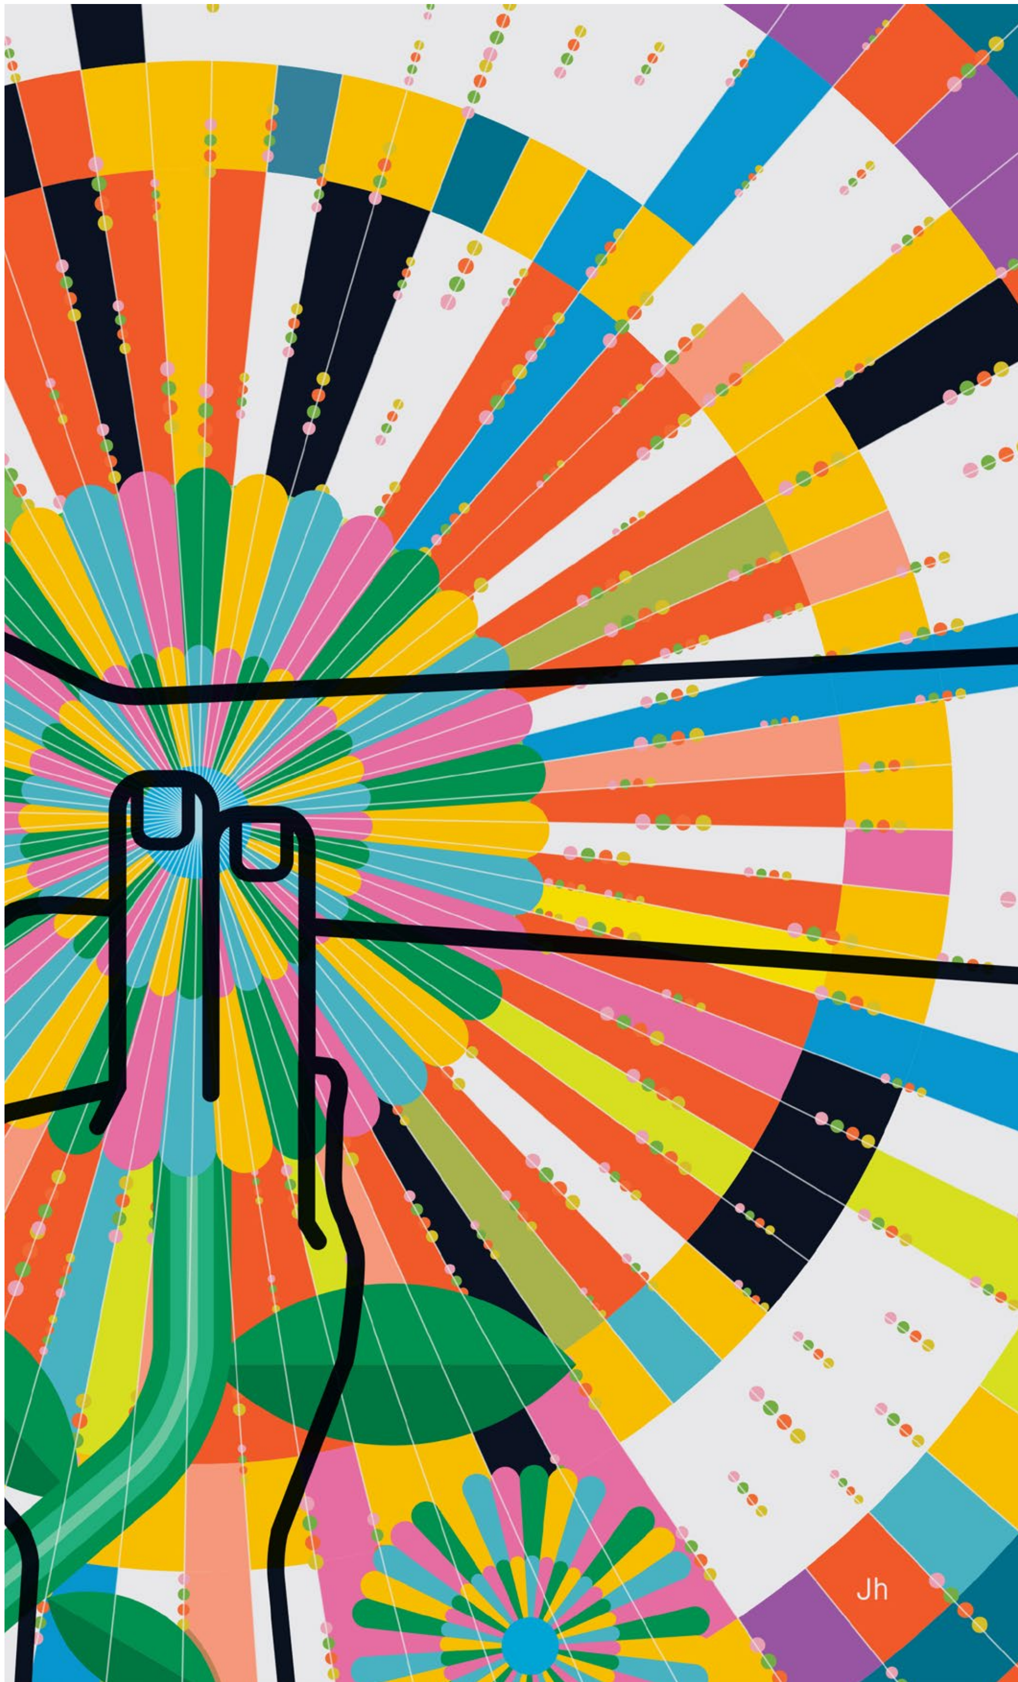

STANFORD MEDICINE

Issue 3 / 2023

TEENS AND SCREENS

HOW TO SAFEGUARD YOUNG PEOPLE'S WELL-BEING ON SOCIAL MEDIA

Social media can be a mixed bag for teenagers. Young people can benefit from online friendships, especially if they feel isolated in real life — for reasons ranging from racial or sexual identity to learning disabilities and geographic isolation.

But research also suggests mental health risks for young social media users — such as increased anxiety, depression, sleep disruption, and struggles with body image and self-esteem. Exposure to cyberbullying and hate-based content are also threats.

Teens understand this pro-and-con balance better than many adults think, according to Stanford Medicine's Vicki Harrison, program director at the Stanford Center for Youth Mental Health and Wellbeing. She has advice on how parents and teens can work together toward healthier online lives.

Vicki Harrison, program director of the Stanford Center for Youth Mental Health and Wellbeing, advocates for more teen involvement in the design of social media platforms they use.

A recent U.S. surgeon general's advisory says social media platforms should embed safeguards for young people. How could that happen?

I would like to see more involvement of young people in the design of social media platforms they use. Our team, in collaboration with Stanford University's design school, recently held a workshop on social media and youth mental health with local high school and college students.

We discussed how several states are considering legislation to protect young people on social media platforms by verifying their ages, having stricter privacy settings enabled by default, minimizing tracking and more.

The big question is how to put such laws into practice. For instance, how should age verification happen, and will the voices of young people be considered? Do they want to have someone scan their faces to verify their ages, or do they want a parent or third-party vendor to verify that? What are the privacy implications of these options?

Participants discussed the tension between protecting young users and social media platforms' freedom to share a wide variety of speech. We heard from young people who feel that, since algorithms already manipulate social media content to create a curated experience, their health and safety should be prioritized over platforms' free speech or profit.

What conversations should parents have with their children about social media use?

I advise parents to delay access to social media as long as they feel they can. The longer your brain develops and your life experiences stack up before you begin using social media, the better outcomes you'll have. I also advise that teens gain access to social media gradually.

It's a good idea to develop a family media plan addressing key questions for teens: Are there rules for when you'll access your device? Who is paying for it? What should you not do on your device? Where can you go for help if you get into trouble? Having those conversations up front and keeping the conversations going once they have access to social media is important.

What myths do you encounter about young people's online lives?

That adolescents want to explore the deep corners of the internet, no matter how horrible. The reality is that they generally want a good experience online. Most turn to social media to laugh, play, connect and learn.

Other teens regard social media as a necessary evil that they join because all their friends are there. They say they lose social currency if they aren't participating, but they don't like a lot of what they encounter.

I have heard from girls that they don't want to see so much content about fitness and losing weight. They don't want that pressure; they just want to have fun and connect with their friends. They are open to adults putting in place commonsense, supportive guardrails. — ERIN DIGITALE

LESLIE WILLIAMSON

VOLUME 40 NUMBER 3
ISSUE 3/2023

STANFORD MEDICINE

SPECIAL REPORT

AI explodes

Taking the pulse of artificial intelligence in medicine

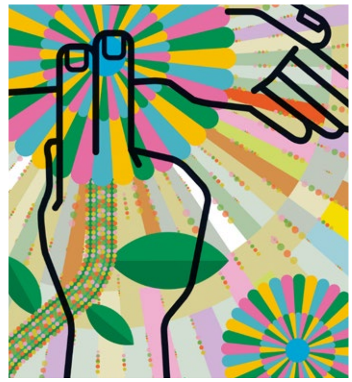

Ways to use AI in health care and research page 6

PLUS

Why genetic analysis early in breast cancer might lead to fewer recurrences page 38

Ingestible capsules plumb the mysteries of the small intestine page 3

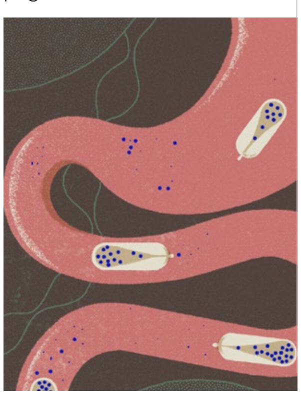

COVER: JOHN HERSEY

DEPARTMENTS

letter from the dean

Imagine a future where your doctor has an AI medical assistant by their side – distilling, in seconds, a world's worth of medical research into a personalized treatment plan for you.

What if, at the click of a button, a researcher could design a custom molecule with the potential to treat a previously untreatable disease? With artificial intelligence's rapid emergence, we are barreling toward this reality.

Academic medical centers around the world, including Stanford Medicine, have begun investigating how AI, including large language models such as ChatGPT, can help us improve patient care, reduce clinician workload, better understand complex biological systems and accelerate drug discovery.

As we embrace this future, we must do so with our eyes wide open. We already have plenty of examples of how AI has fallen short in biomedical research — typically because of biased or otherwise faulty data — and there will no doubt be further unforeseen consequences. As with any powerful new tool, we must not only develop the knowledge and skills to employ it effectively but also invest in shaping rules to guide its safe and responsible use.

Recognizing the urgent need to define AI's place in health and medicine, Stanford Medicine and the Stanford Institute for Human-Centered Artificial Intelligence recently launched Responsible AI for Safe and Equitable Health, or RAISE-Health. The goal of this trailblazing initiative is to guide the principled use of AI across biomedical research, education and patient care.

Central to RAISE-Health's mission is ensuring that this new technology does not worsen current health inequalities but rather helps eliminate them. We are also creating a platform we hope will serve as a go-to resource to enable academic, government and industry leaders to make informed decisions. We know no system will be free of imperfection, but to say our current environment has room for improvement would be an understatement. By moving forward with intention and purpose, Stanford Medicine can ensure that successes far outnumber any stumbles along the way.

I'm proud that Stanford is asserting this leadership position. Our expertise in computer science and artificial intelligence extends back to the infancy of these fields. Our world-class biomedical and bioinformatics faculty drive groundbreaking discoveries daily. And with Silicon Valley partners, we have access to the most potent innovation hub on the planet. By bringing together decision makers, experts and diverse voices, we are uniquely positioned to define ethical standards and safeguards for AI in medicine.

As we stand on the cusp of this revolution and imagine how our lives and roles will change, I believe that AI's impact will rival that of some of the most transformative innovations of human history, including the printing press and the internet. Progress will not be linear, but initiatives such as RAISE-Health will be critical in establishing best practices to secure a healthier and more equitable future for people around the world.

Sincerely,

Lloyd Minor, MD

Carl and Elizabeth Naumann Dean of Stanford School of Medicine

Vice President for Medical Affairs at Stanford University

Professor of Otolaryngology-Head & Neck Surgery

PHOTOGRAPH BY CAROLYN FONG

upfront

Squirming solution

CHILDREN WHO underwent radiation treatment for cancer were less likely to need anesthesia if they were watching videos during the procedure, a study led by Stanford Medicine researchers has found.

With video distraction, 78% of children in the study could hold still through at least one 10- to 30-minute radiotherapy session without anesthesia. Prior studies found that less than half of the children could tolerate radiotherapy without anesthesia.

The study was published March 29 in the *International Journal of Radiation Oncology, Biology, Physics*.

Most young kids who need radiotherapy get general anesthesia to keep them still during treatment, which allows the radiation beams to be aimed precisely at their tumors.

"If we can get them engrossed in paying attention to something, such as a video they enjoy, that really helps," said Susan Hiniker, MD, assistant professor of radiation oncology and a senior author of the study.

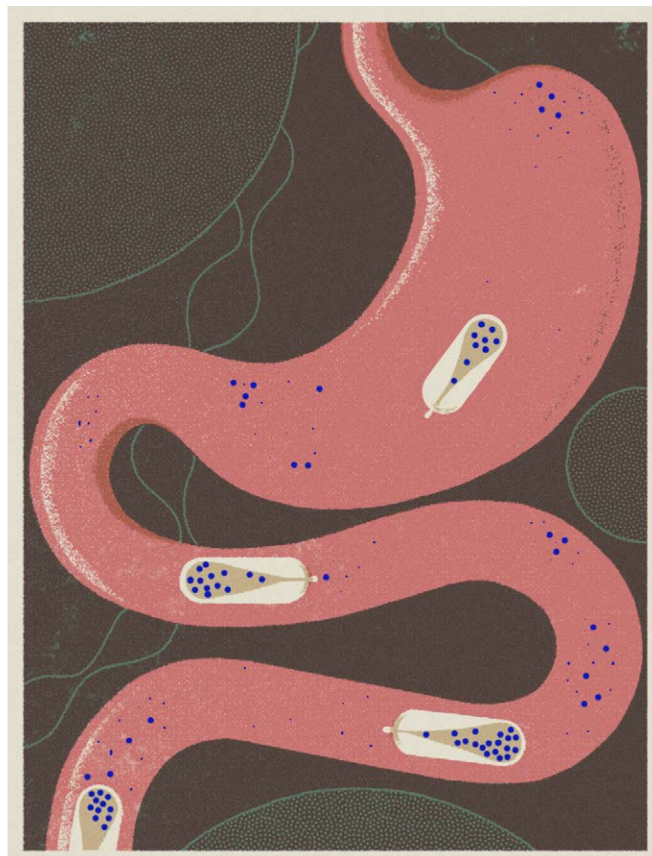

Fantastic (intestinal) voyage

STANFORD MEDICINE researchers and their collaborators have developed an ingestible device that can document the diversity of microorganisms, viruses, proteins and bile salts in the small intestine.

The proof-of-concept results were published May 10 in *Nature*.

"Samples from current tools don't fully represent what's going on inside of us. But it's all we've had — until now," said KC Huang, PhD, a professor of bioengineering and of microbiology and immunology.

Huang shares senior authorship with David Relman, MD, a professor of medicine and of microbiology and immunology.

Other ingestible devices have been invented to sample gut microbiota, but they've been loaded with electronics, limiting their widespread use due to manufacturing complexity and cost.

The new device is a hollow capsule enclosing a bladder that sucks in a sample. Various versions of the capsule open at different pH levels. Since the acidity of the small intestine's contents decreases along the tract, a set of capsules is able to sample an array of the organ's microenvironments.

"I think the magic of our capsule is that it is going to transform how people think about the gut microbiota's relationship to diseases," Huang said.

> 'I THINK THE MAGIC OF OUR CAPSULE IS THAT IT IS GOING TO TRANSFORM HOW PEOPLE THINK ABOUT THE GUT MICROBIOTA'S RELATIONSHIP TO DISEASES.'

ILLUSTRATIONS BY JUAN BERNABEU

Healing tips from worms

THE PLANARIAN HAS A NEAT TRICK: If it loses its head, it'll grow a new one. Of course, it's a flatworm, so practicing the trick yourself is discouraged. Still, understanding how the worm does it could lead to medical advances, such as new cancer treatments and ways to regenerate tissue after injury.

Bo Wang, PhD, assistant professor of bioengineering, was intrigued by the fact that in planarians and many other animals, including mice, zebrafish and axolotls, wounds in one part of the body trigger responses in distant tissues.

"But whether or not those responses actually have any function has been unclear," Wang said.

He and his team wanted to understand how these responses coordinate, so they conducted a series of experiments in planarians. Cells in animals from worms to humans use a chain of proteins called the extracellular signal-related kinase, or ERK, pathway to communicate with each other. If tissue is injured, the nearest cells pass along that information to neighboring cells, which then tell their neighbors, and so on, leading to a signal wave propagating in the animal's body.

To find out whether cells far from an injury engage in the healing process through this pathway, Yuhang Fan, a graduate student in the Wang lab, blocked ERK signals from spreading in a planarian. Then he cut off its head.

Normally, a flatworm's head quickly regrows after removal. But in this case, the head never regrew. Then, Fan removed a planarian's head and its tail. Surprisingly, both regrew. Taken together, these findings indicate regeneration required wound responses in tissues throughout the animal's body.

"This implies there's kind of a global body voting system that says, 'OK, now we should grow something,' and everybody has to agree within a short time after injury," Wang said. And even the cells furthest away get a vote.

As the researchers tracked ERK-signaling waves spreading throughout planarians' bodies, they noted that hundreds of genes were turned on and off. Although quite distantly related to planarians, humans share many of those genes.

"This really gives us an entryway to go after those genes," Wang said. "It could allow us to figure out how animals regenerate while managing the risk of uncontrolled cancerous growth."

'IT'S ALMOST AS IF YOU'D ALREADY DECIDED HOW YOU WERE GOING TO FEEL, AND THEN EVERYTHING YOU WERE SENSING WAS FILTERED THROUGH THAT.'

Reversing depression

A STUDY LED BY Stanford Medicine scientists has revealed that transcranial magnetic stimulation treats depression by correcting the abnormal flow of brain signals.

Previously, it was a mystery why the therapy, in which magnetic pulses are used to stimulate neurons, was effective in people with the mental illness. In the brain, the anterior insula, a region that integrates bodily sensations, typically sends signals to a region that governs emotions, the anterior cingulate cortex.

"You could think of it as the anterior cingulate cortex receiving this information about the body — like heart rate or temperature — then deciding how to feel on the basis of these signals," said Anish Mitra, MD, PhD, a postdoctoral fellow in psychiatry and behavioral sciences and lead author of the study published May 15 in the *Proceedings of the National Academy of Sciences.*

Researchers analyzed brain-imaging data from 33 study participants with depression. In three-quarters of them, the typical activity flow was reversed — the anterior cingulate cortex sent signals to the anterior insula. The more severe the depression, the higher the proportion of signals that traveled the wrong way.

"It's almost as if you'd already decided how you were going to feel, and then everything you were sensing was filtered through that," Mitra said. "The mood has become primary."

In patients treated with Stanford neuromodulation therapy, a form of transcranial magnetic stimulation — pioneered by Nolan Williams, MD, associate professor of psychiatry and behavioral sciences — the neural activity flow shifted to the normal direction within a week, coinciding with a lifting of their depression.

Poor memory linked to autism

CHILDREN WITH autism not only have difficulty remembering faces but also recalling other kinds of information, a recent study reported.

Stanford Medicine scientists conducted the research and reported the findings July 10 in *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging.*

Social challenges are a core feature of autism, and it's possible that memory impairments significantly contribute to the ability to engage socially, researchers said. "Social cognition cannot occur without reliable memory," said senior author Vinod Menon, PhD, the Rachael L. and Walter F. Nichols, MD, Professor and a professor of psychiatry and behavioral sciences.

Past studies have shown that children with autism have trouble remembering faces. Some small early studies also suggested that children with autism have broader memory difficulties.

In the new study, high-functioning children with autism and normal IQs scored lower than typically developing children on tests of immediate and delayed verbal recall, immediate visual recall, and delayed verbal recognition.

Brain scans of the children with autism showed that distinct brain networks drove different types of memory difficulty.

"The findings suggest that general and face-memory challenges have two underlying sources in the brain that contribute to a broader profile of memory impairments in autism," Menon said.

The beat goes on

FOR THE FIRST TIME, a heart from an organ donor who died of cardiac arrest was restarted and then transplanted while it was beating. Initially performed by Joseph Woo, MD, professor and chair of cardiothoracic surgery, and his team in October 2022, the operation has since been used in adult and pediatric patients more than a dozen times by Stanford Medicine surgeons.

Those who have experienced brain death — not cardiac arrest — have long been the source of most heart transplants. That's because brain death donors are kept on life support, keeping their heart beating, which helps maintain the organ's health.

But with demand outpacing supply, the medical world has been pushed to seek new approaches. In the U.S. about 3,500 people await a heart transplant.

Recent technological advances have allowed for more success with hearts from donors who died by what's known as cardiac or circulatory death, in which the heart has already stopped once, either naturally or because life support was discontinued.

Such procedures increase the number of hearts available for transplant, but outcomes for the recipients are poorer. These hearts have traditionally been stopped twice — first at death, then immediately before transplantation, after spending time hooked up to a device that perfuses them with oxygenated blood while outside of the body.

"Stopping the heart a second time, just before transplanting, induces more injury," said Woo, the senior author of a study describing the beating heart procedure that published in the *Journal of Thoracic and Cardiovascular Surgery Techniques* in March. "I asked, 'Why can't we sew it in while it is still beating?'"

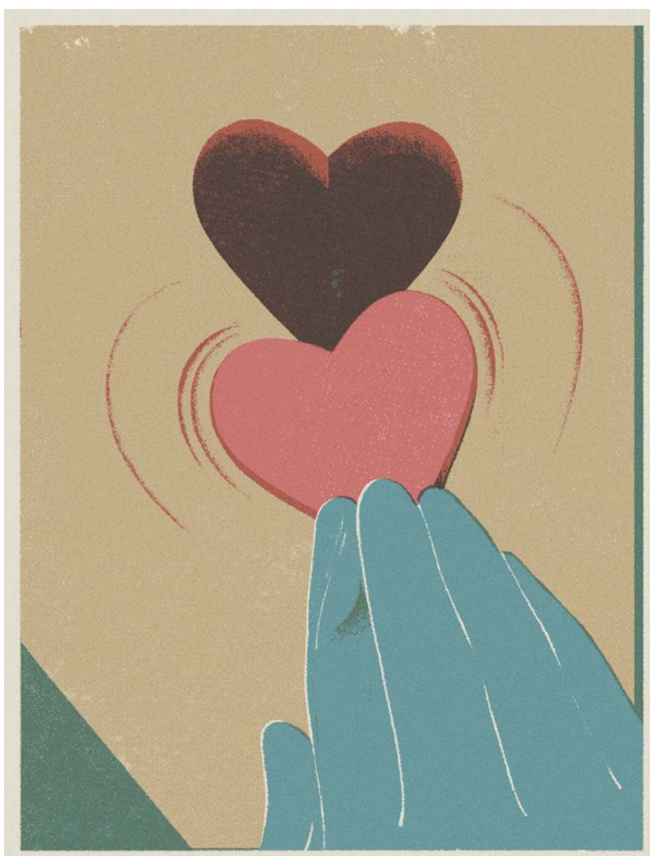

Healthy fat's longevity link

RESEARCHERS studying tiny transparent worms have identified a cellular connection between lifespan and healthy fats known as monounsaturated fatty acids.

The findings reported May 1 in *Nature Cell Biology* hint at the complex relationship between diet, fats and longevity.

"Fats are generally thought to be detrimental to health," said Anne Brunet, PhD, the study's senior author and the Michele and Timothy Barakett Endowed Professor. "But some studies have shown that specific types of fats, or lipids, can be beneficial."

The researchers learned that roundworms fed foods rich in the fat building block oleic acid, a key component of the Mediterranean diet, lived about 35% longer than those on standard rations.

They also found that oleic acid raises the number of two cellular structures — lipid droplets and peroxisomes — and protects cells' membranes from damage by oxidation. Surprisingly, lipid droplets served as a de facto crystal ball for predicting the number of days each animal would live.

"The number of lipid droplets in individual worms tells me that animal's remaining lifespan," said research scientist Katharina Papsdorf, PhD, the study's lead author. "The worms with greater numbers of lipid droplets live longer than those with fewer droplets."

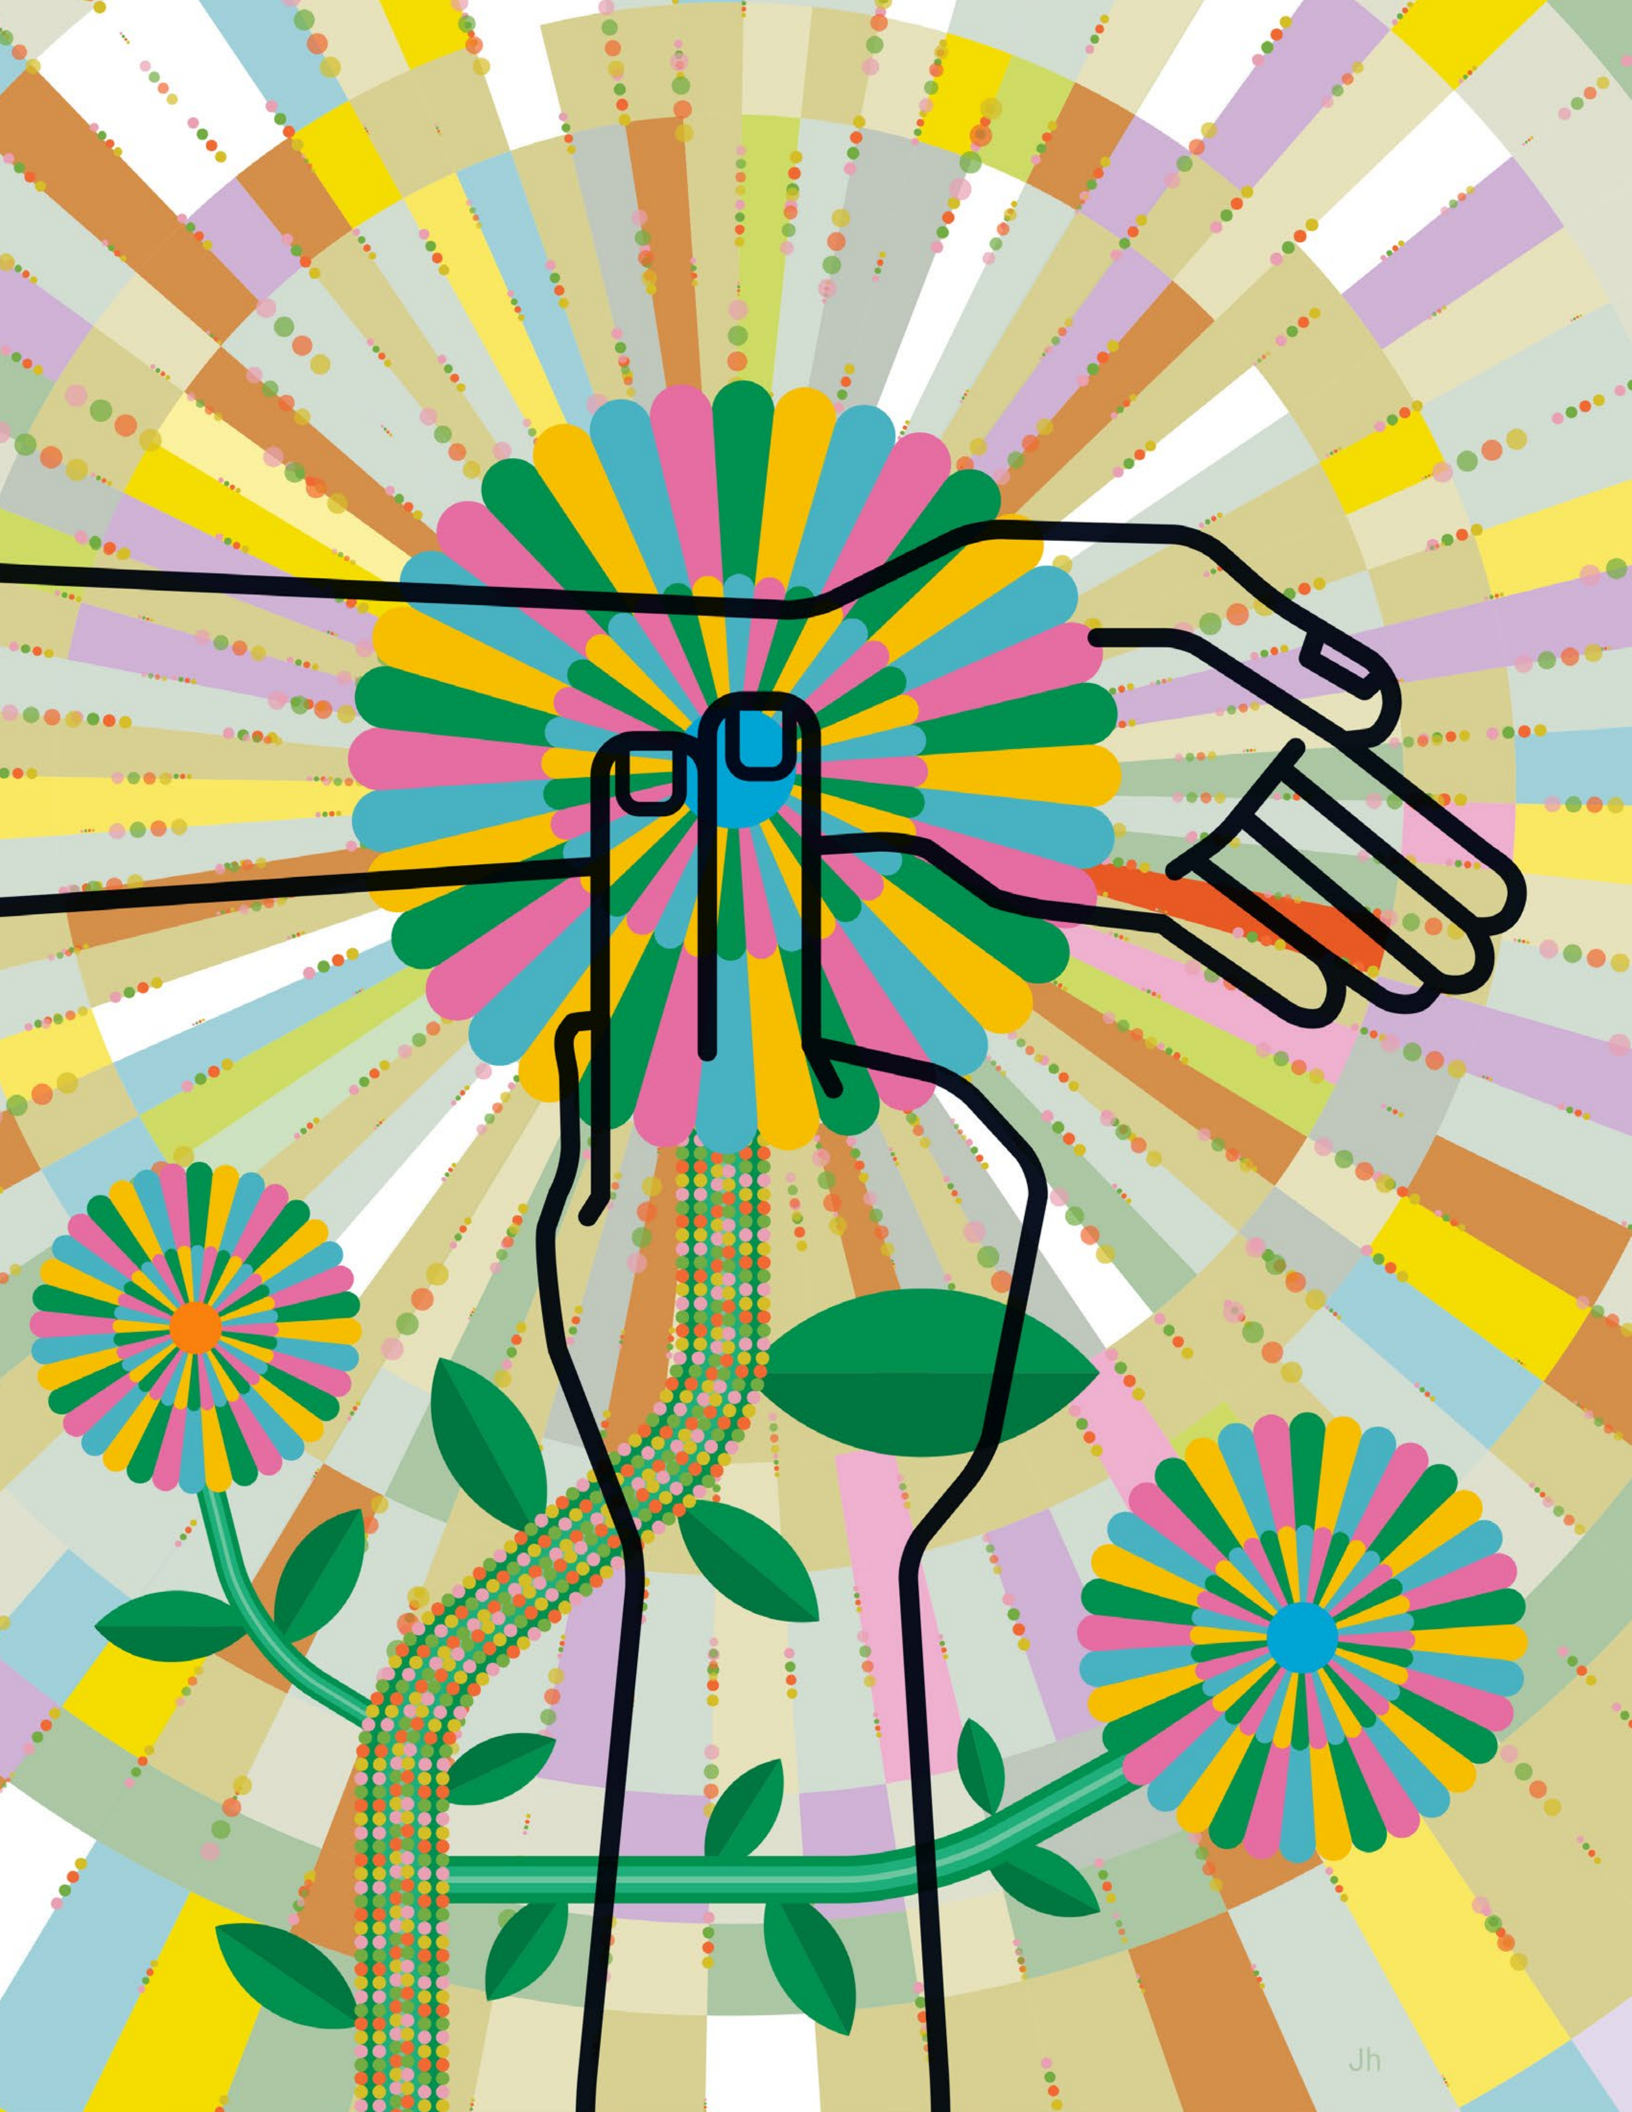
Jh

MEDICINE'S AI BOOM

The Stanford impact

Introduction by Mark Conley **Vignettes by Anna Marie Yanny** **Illustration by John Hersey** **Portraits by Misha Gravenor**

Outside the Stanford Health Care office of data scientist Nigam Shah hangs an antiquated memento from the original wave of artificial intelligence hype of the mid-1960s. Yes, the '*60s*. ● Against a backdrop of free love and Vietnam War acrimony, the first AI wave washed over Stanford University and numerous other academic institutions with hardly a ripple felt by a distracted outside world. ● Its arrival in the early days of Silicon Valley coincided with that of two pioneering computer scientists who came to help Stanford launch the country's first "super computer" for AI in medicine. Housed at the Lane Library, it was called — as the stained-glass sign in shades of oceanic blue now hanging in Shah's office reads — SUMEX-AIM, short for Stanford University Medical Experimental Computer for Artificial Intelligence in Medicine.

The term "artificial intelligence" had been coined a decade earlier by mathematics and computer science professor John McCarthy, PhD, who came west from MIT to launch the influential Stanford Artificial Intelligence Laboratory in 1963.

In other words, even with all the buzz surrounding the current AI explosion, it's really nothing new — especially at Stanford.

"If you poll the 1,200 faculty in the School of Medicine, I'd be surprised if more than 10% know about any of Stanford's AI history," said Shah, MBBS, PhD, professor of medicine and of biomedical data science and chief data scientist for Stanford Health Care. "A lot of people think this right now is the first AI hype cycle."

It helps explain why Shah, who has been terabytes-deep in machine learning and neural networks for decades, takes his duties as an AI historian and pragmatist seriously. He wants newbies just tuning in to understand the context when deciphering what has deep relevance and what is just another round of futuristic noise.

Shah describes this as a moment of both high frenzy and immense opportunity, with a venture-capital-fueled rush to deliver applications with lasting value — a goal he estimates only 5% to 10% of the applications are hitting in today's influencer-inspired culture that seeks "breakthroughs every 24 hours." Through that haze of ambition, real-world innovations in medicine are emerging. They are just more subtle than sexy; more incremental than game-changing.

As Stanford Health Care's AI vetter-in-chief, Shah is watching closely. If algorithms are designed right and serve a useful purpose, even hype busters like him are willing to buy into this wave of anticipation that has washed over every sector of society. Expectations are especially high at a place like Stanford Medicine, where the interests of industry and academia synthesize in a way that they do at few other medical institutions.

"If we utilize this moment of attention, it's quite an opportunity," Shah said, adding that it will take the best ideas being tested at the country's 600 health care systems — no small task. "If we do that," he added, "we will have pulled off a national experiment at a scale that no government agency or single company could have done. That is immensely exciting."

It's a historic opportunity, and it raises questions about how to use AI in medicine responsibly, how to set realistic expectations for its potential and what part the humans behind the algorithms will play. Here are a few of the most pressing questions.

FOR A GLOSSARY OF AI TERMINOLOGY, SEE 'ABCs OF AI' ON PAGE 23.

What is Stanford Medicine doing to ensure AI is used responsibly in research and health care?

FEARS ABOUT AI ARE REAL, particularly in the sensitive world of health care. Could these new algorithms compound existing challenges such as bias in how people are treated based on their race and loss of privacy due to health data breaches? Could they ratchet up distrust in the health care system and those who provide care?

In June, a collaboration between the Stanford Institute for Human-Centered Artificial Intelligence, or HAI, and the School of Medicine resulted in the RAISE-Health initiative, whose mission is to "guide the responsible use of AI across biomedical research, education, and patient care."

Co-sponsored by Lloyd Minor, MD, dean of the School of Medicine and vice president for medical affairs at Stanford University, and Fei-Fei Li, PhD, co-director of HAI and professor of computer science, the initiative brings together a diverse set of voices from across the Stanford community — ethicists, engineers and social scientists.

Feeding that knowledge pool are researchers like Tina Hernandez-Boussard, PhD, associate dean of research and professor of medicine, of biomedical data science and of surgery, whose work aims to ensure that diverse populations receive equitable resources and care. And that of Sanmi Koyejo, PhD, a RAISE-Health co-leader, whose work centers on fairness and the detection of potential bias in AI-aided analysis of medical imaging.

"AI is poised to revolutionize biomedicine, but unlocking its potential is intrinsically tied with its responsible use," Minor said. "We have to act with urgency to ensure that this technology advances in everyone's best interests."

That will involve sound government regulation, something Sherri Rose, PhD, a professor of health policy, is shaping. Rose's research has focused on making sure the interests of marginalized populations are considered in the rush of AI research.

And what about how we're training a new generation — perhaps the first real AI generation — on how to deal with all of the new issues that AI in medicine is raising? There's a fellowship for that. Through a program led by the Stanford Center for Biomedical Ethics, with funding from pharmaceutical company GSK, three postdoctoral fellows are exploring the ethical, legal and social considerations arising from the use of artificial intelligence in the pharma industry, from early-stage drug discovery efforts to its use in doctors' offices.

How are we defining AI?

BETWEEN DEEP NEURAL NETWORKS, machine learning, large language models and generative versus nongenerative AI — just for starters — the jargon of artificial intelligence can get fuzzy fast. Even the term "artificial intelligence" has multiple meanings.

"It depends who you're talking to," Shah said. "It can be in the mind of the beholder."

Its original meaning is the imitation by computers of human intelligence. That's a tall order — and might never come to pass. More often people use the term to refer to computers that can accomplish complex tasks, but there's no clear divide between a plain old algorithm doing run-of-the-mill computing and an AI program. Is a pocket calculator AI? Most would say no. But the face recognition feature on a smartphone would make the grade.

Shah's "AI 101" list of definitions looks like this:

- Machine learning is the basis for building algorithms or models by feeding a computer certain datasets. It is where all AI problem-solving begins.
- A neural network is the most complex form of machine learning, using all possible mathematical inputs with no constraints — and with multiple layers, it can become a deep neural network, producing what is known as deep learning.
- If you can take that algorithm or model produced by machine learning and generate data with it, it's called generative AI. Ironically, ChatGPT, the buzz-producing model du jour, "makes stuff up," Shah said. "That's what it's designed to do."
- Large language models, such as the internet-fed GPT-4.0 and Bard, are the rocket fuel bringing generative AI to life.
- Shah is clearest on one point: Even if generative AI is responsible for all the attention, it will not factor prominently into medicine and health care any time soon because of the concerns around accuracy and trust.

Nor, he said, will 90% of the algorithms being produced today.

"Ten years from now, we'll be immensely grateful for the 10% that panned out and changed the science, the practice, or the delivery of care in medicine," he said.

In what ways is Stanford Medicine taking the AI lead?

THE FORMATION OF the Stanford Center for Artificial Intelligence in Medicine and Imaging (known as AIMI) in 2018, which brought together 50 faculty from 20 departments, spurred the wide range of interdisciplinary team science being done at Stanford Medicine in machine learning today.

Shah was among the first to embrace such collaboration. As Stanford Health Care's chief data scientist, he is on the front lines of bringing AI from the research lab to the clinic. His team is looking for the immediate winners that could serve an important purpose for both doctors and patients.

At Stanford Health Care, the early focus with generative AI has been on how large language models can be used to communicate with patients. Can, for instance, AI support tasks like note taking so doctors can be more engaged with their patients? Can AI help organize patient care records — particularly for multiple providers — more efficiently?

"Paradoxically, I think AI will help take the computer out of the room and allow the two humans to make a closer human bond," said Euan Ashley, MBChB, DPhil, associate dean and professor of medicine, of genetics and of biomedical data science. "And I think medicine will be better for that."

Patricia Garcia, MD, a clinical associate professor of gastroenterology and hepatology, is leading a pilot program in which a generative AI tool based on ChatGPT automatically drafts responses to patients' medical advice requests for clinicians to review, edit and send.

Russ Altman, MD, PhD, a professor of bioengineering, of genetics, of medicine and of biomedical data science, and an associate director of HAI and one of four co-leaders of the RAISE-Health initiative, has focused largely on AI's applications for drug discovery and how those drugs will work on patients — along with the ethical implications involved.

Shah's biggest concern centers on how the froth this AI buzz has generated is affecting the iteration process. A lot of what is being created, he said, is formed backward. Rather than follow what is known as a biodesign process — where need dictates solutions — the AI buzz has many researchers reaching blindly for the next great thing. "There's a lot of poorly designed hammers looking for nails," Shah said.

In what areas is AI showing the most immediate promise?

IF STANFORD MEDICINE'S AI imprint to date required a single-word summation, it would be imaging. Of the 500-plus AI algorithms approved by the FDA, 75% are radiology-focused and 85% are imaging-focused. At Stanford Medicine, AI innovation skews heavily toward imaging as well.

In radiology, imaging technologies such as X-ray and MRI are used to diagnose patients, and the field has produced one of the few robust and consistent datasets in medicine. The human genome is another. At their intersection lies a perfect example of the opportunity for discovery.

"AI can be, in some ways, superhuman because of its ability to link disparate data sources," said Curtis Langlotz, MD, PhD, professor of radiology, of medicine and of biomedical data science as well as the director of the Center for Artificial Intelligence in Medicine and Imaging. "It can take genomic information and imaging information and potentially find linkages that that humans aren't able to make."

Langlotz, who's also a co-leader of RAISE-Health and associate director of HAI, has been at the forefront of AI and imaging for many years. His lab feeds medical images and clinical notes into deep neural networks with the goal of detecting disease and eliminating diagnostic errors that could occur with only a doctor's assessment.

As Langlotz puts it: "Cross-correlating massive datasets is not well suited to humans' information-processing capability."

Ashley, who directs Stanford Medicine's Center for Inherited Cardiovascular Disease, has been pushing algorithms with similar success at the intersection of genetics and cardiology in an effort to improve the detection of cardiac risk. He said there is little doubt that highly trained computers can synthesize medical data in ways that humans cannot. But it's now a matter of taking early methodology successes from the lab into clinical trials and proving they work in humans.

"We're at an interesting moment when we've demonstrated there is power — that computers can predict things that a human cannot," he said. "Now, how do we get it into medical practice?"

'WE'RE AT AN INTERESTING MOMENT WHEN WE'VE DEMONSTRATED THERE IS POWER — THAT COMPUTERS CAN PREDICT THINGS THAT A HUMAN CANNOT.'

What are some of Stanford Medicine's other AI focuses?

Christina Curtis, PhD, RZ Cao Professor of medicine, of genetics and of biomedical data science, is analyzing the molecular profiles of tumor samples and integrating routine pathology images to advance the standard of care for breast cancer patients. Adding detailed genomic data to the mix, clinicians might one day be able to pinpoint the best treatment for each patient.

"Currently, most cancer patients undergo sequencing only once they've developed treatment-resistant metastatic disease," Curtis said. "There is a missed opportunity to have such information earlier in the disease course, at the time of initial diagnosis, both to compare a given patient to other similar patients and to monitor how the disease changes over time. This could enable more precise and anticipatory care."

Sylvia Plevritis, PhD, a professor of biomedical data science and of radiology, is also one of the leaders of the RAISE-Health initiative. She developed a systems biology cancer research program that bridges genomics, biocomputation, imaging and population sciences to decipher properties of cancer progression. "Today, AI is completely changing the way we connect the dots between basic science research and clinical care," said Plevritis, who is in the Stanford Medicine leadership group focusing on what this new path means for medical education.

"We're having to rethink what scholarship and creativity are when we have tools that can write for us," she said.

It's one of the many unanswered questions that make this moment equal parts precarious and exhilarating. Here is a closer look at the work of some of the many Stanford Medicine humans trying to navigate the AI maze in search of answers.

FORCE MULTIPLIERS

AI HELPS PEDIATRICIANS CHECK HEART HEALTH

Speedier, easier heart-pumping assessment for children

AN IMPORTANT INDICATOR of the heart's function is its ability to pump blood to the body — an action typically powered by the left ventricle. But estimating the amount of blood pumped with each heartbeat is time-consuming, and measurements can vary between cardiologists. This assessment has been automated for adults but not pediatric patients.

So, pediatric cardiologist Charitha Reddy, MD, teamed up with engineers and computer scientists to develop a model that automatically estimates the left ventricle's function in children


CHARITHA REDDY

with accuracy and reliability. Medical decisions for children that rely on a doctor's assessment of the heart's ability to pump, such as determining safe chemotherapy doses, can benefit from models tailored to children, she said.

Reddy, a clinical assistant professor of pediatrics, helped collect heart ultrasound videos and annotate images from 1,958 pediatric patients seen at Lucile Packard Children's Hospital Stanford. The model she helped develop analyzed more than 4,000 video clips of hearts and generated estimates of the left ventricle's function with high accuracy. Its assessments of heart-pumping ability were speedier and more consistent than doctors', said Reddy, the lead author of a study of the model published in February in the *Journal of the American Society of Echocardiography*.

"It performed essentially as well as a human doing the same measurements, while being less variable," Reddy said. "The model could improve efficiency with this task or serve as an 'expert' when there isn't one available."

Clinicians could use the model to assess heart-pumping ability with confidence that the measurement is within about 3% of what they would measure themselves, Reddy said.

The algorithm could offer a template for AI models to assess the function of the right ventricle and the pumping ability of fetal hearts and hearts with structural abnormalities.

The model needs further testing before it's put to use, but Reddy hopes it will someday help noncardiologists, such as parents of children with heart problems who live far from a hospital, screen for heart-pumping weakness. She also thinks it could improve care in rural areas with fewer cardiologists.

ROXANA DANESHJOU

BETTER PHOTOS OF SKIN FOR TELEHEALTH VISITS

In the age of virtual doctor appointments, this app could improve patient photos and expedite treatment

WHEN THE COVID-19 pandemic hit in 2020, Roxana Daneshjou, MD, PhD, was a resident in Stanford Medicine's dermatology clinic, where she helped provide screenings for skin conditions, such as shingles, eczema and suspicious moles. But as the clinic was forced to pivot from in-person visits to telehealth, doctors had to rely on photos taken by patients, which were often hard to interpret. Daneshjou and her colleagues spent hours combing through blurry photos with bad lighting.

"I was reviewing these photos and thought, 'I think we can develop an algorithm to assess this automatically,'" said Daneshjou, who was also researching how to apply artificial intelligence in health care. "Maybe it could help patients submit clinically useful photos."

To that end, Daneshjou and a team of researchers collected images of skin conditions depicting a variety of skin tones and used them to train an algorithm to identify low-quality photos — and recommend ways to fix them. Funded by Stanford's Catalyst program, which supports medical innovations on campus, Daneshjou developed the algorithm into a web app called TrueImage that patients can access from a smartphone or tablet. The idea is that patients will follow the app's prompts to snap a photo that's good enough for their doctor. TrueImage rejects low-quality images and tells patients to move to a brighter room, zoom in on their lesion or sharpen the focus. Currently, there's a long wait for new dermatology appointments nationally, which the app could expedite by helping doctors quickly procure quality images, said Daneshjou, who is now an instructor of biomedical data science and dermatology.

But algorithms like hers matter only if humans actually use them and change their behavior, Daneshjou said. In a 2021 pilot study run in the clinic, she found TrueImage reduced the number of patients who submitted poor images by 68%. Now, she's running a larger clinical trial to probe the app's efficacy when patients use it at home.

If the trial confirms the app's ability to improve photos that patients submit, the team will launch it in Stanford Medicine clinics. Daneshjou hopes the app can also help doctors outside of dermatology, like primary care providers, collect quality images as they screen for skin conditions.

CHEST SCANS CHANGE PATIENTS' MINDS

AI analyzed repurposed chest CT images to identify calcium buildup in arteries, which encouraged patients to make lifestyle changes

CARDIOLOGIST Alexander Sandhu has many patients who could benefit from a CT scan of their heart to detect the buildup of calcium inside their coronary arteries. The plaque is the strongest risk predictor of heart attacks. However, most patients opt out of this telling scan, often because their insurance doesn't cover it. But a few years ago, Sandhu noticed a happy coincidence.

"Quite frequently, patients have already had a chest CT done for some reason totally unrelated to their heart," said Sandhu, MD, assistant professor of medicine in the division of cardiovascular medicine. He wondered if artificial intelligence could discern valuable heart information in chest scans previously carried out for any reason — for example, to screen for lung cancer. "I thought, there's this incredible opportunity to repurpose those scans and provide that information to patients and their clinicians."

Along with his mentor, David Maron, MD, director of preventive cardiology at the School of Medicine, Sandhu helped design a deep learning algorithm that could assess the amount of calcium in patients' coronary arteries from a chest CT scan with about as much accuracy as a scan ordered for that specific purpose. In 2022, the team tested the algorithm in 173 patients, most of whom were at high risk for heart disease but were not taking statin medications, which are known to decrease the risk of heart attack and stroke.

They found that when they notified the patients and their primary care physicians about the calcium detected and showed them images of the white deposits in the patients' arteries, 51% of them started a statin medication within six months. That's about seven times the statin prescription rate the team observed in a similar group they did not notify, Sandhu wrote in an analysis published November 2022 in the journal *Circulation*. (At the end of the study, the researchers notified all patients with coronary calcium.)

"These are usually patients who have no symptoms, and this provides motivation to make lifestyle changes and take statin medications," Maron said.

In 2023, their research was awarded the James T. Willerson Award in Clinical Science and a Hearst Health Prize. The notification system is currently in use at about 50 U.S. hospitals, Sandhu said.

Sandhu, Maron and their research collaborators are now planning a clinical trial to test whether the system prevents heart attacks and strokes.

ALEXANDER SANDHU

MAAME YAA A.B. YIADOM

COULD AI RIVAL AN EYE SPECIALIST?

An AI model could predict whether patients will need eye surgery to prevent vision loss

WHEN IT COMES TO GLAUCOMA, patients' eyes are fine — until they're not. Many with this eye disease remain asymptomatic for years. Yet others with the condition rapidly progress toward irreversible blindness — and may need surgery to prevent it.

"It's not easy to figure out which patients are which," said glaucoma specialist Sophia Wang, MD, an assistant professor of ophthalmology at the School of Medicine.

Wang wondered if AI could help. Recently, language-processing technology has opened doors to swiftly analyze doctors' notes, unearthing a wealth of information — such as family history — to aid in predicting which glaucoma patients will need surgery soon.

Wang and her collaborators trained a two-pronged AI model to "read" doctor's notes and weigh clinical measurements from 3,469 glaucoma patients to understand what health factors — captured by measurements of eye pressure, medications and diagnosis codes — seemed to indicate that they would need surgery in the coming year. The idea is that doctors could use the model to flag riskier patients and start them on aggressive treatments to prevent vision loss, Wang said.

The model can tell doctors a patient's chance of needing surgery within the year with more than 80% accuracy, according to a study published in April in *Frontiers in Medicine*. Doctors would likely want to act quickly if the model predicted a high chance that a patient would need surgery, whereas a patient with a low chance could potentially be safely monitored with less invasive treatment options. While it's not perfect, it gives doctors a sense of how worried they should be about the patient, said Wang, the study's senior author. Surprisingly, she found the model was more accurate than her own assessments.

"If I review 300 notes and try to predict who is going to need surgery, my performance is abysmal, even as a highly trained glaucoma specialist," Wang said. The different, disparate details in health records don't necessarily intuitively add up to needing surgery, Wang said, making the prediction especially difficult for doctors. "The model far outperforms humans in this task."

But before the model can be used in the clinic, Wang and her team want to further refine it by incorporating imaging data. They also plan to test the model on large cohorts of patients who are racially, ethnically and socioeconomically diverse and live in various regions of the United States.

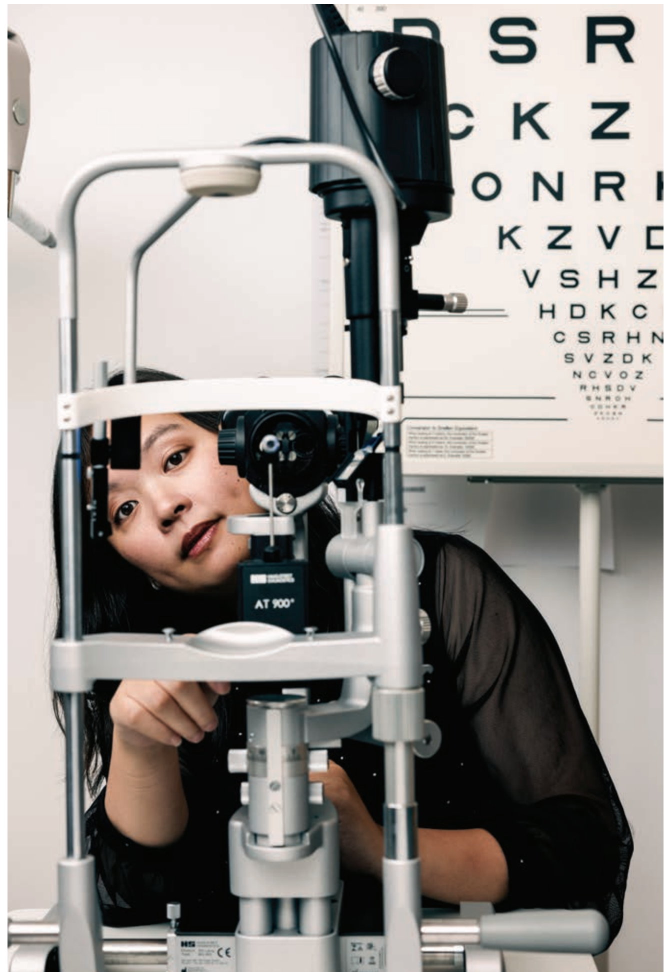

SOPHIA WANG

PATIENT CARE

IMPROVING EQUITY IN HEART ATTACK SCREENING

Together, humans and AI could screen for heart attacks more precisely and equitably

EMERGENCY DEPARTMENT PROTOCOL IS CLEAR. When patients arrive with signs of a heart attack, the ED team has less than 10 minutes to assess them with an electrocardiogram, or EKG, to determine whether a blockage is reducing blood flow to their heart. As the minutes increase, so does a patient's risk of permanent heart damage.

That's a lot of pressure on the staff registering incoming patients to make the right call, said Maame Yaa A. B. Yiadom, MD, associate professor of emergency medicine. She was curious if AI could reduce delays in care. "We intended to build a model that could outperform humans in this screening task," Yiadom said.

Funded by HAI, Yiadom's lab — including a team of statisticians, data scientists, medical informaticists and emergency medicine physicians — created a predictive model that could determine within 10 minutes whether an ED patient should be tested for a heart attack upon arrival with an EKG.

To assess the model's performance, they used it to analyze electronic health records from 279,132 past visits to Stanford Hospital's ED. Using information such as patient age, chief complaint and other routinely collected data the model predicted which patients would go on to be diagnosed with an adverse heart event, meaning they should be screened with an EKG. The model wasn't perfect — it missed capturing 18% of cases, but it outperformed ED staff screening, which typically missed about 27%.

Still, it wasn't the solution Yiadom's team was hoping for.

"Our biggest surprise was that the model was biased when it came to race," she said. In their study, the model was worse at detecting heart attack risk in Native American, Pacific Islander and Black patients than in white patients. The team observed that members of these historically disadvantaged populations tend to have heart attacks at younger ages, and the model generally categorized young patients as low risk, Yiadom said.

"There's something that the humans are doing that is introducing equity into the screening. We don't want to throw that away," Yiadom said.

So, she and her team tested a "fail safe" screening approach, in which patients would get their hearts tested if either the staff or the AI model indicated a need. The human-AI combo method performed better than either alone, missing only 8% of cases while showing little variation between patients of different sexes, races, ethnicities and ages, Yiadom wrote in a study published in June in the journal *Diagnostics*.

She intends to test the combined method in clinics to ensure that it doesn't impede ED processes and that it functions equitably in a real setting. Yiadom is also designing a new AI model that she hopes will be unbiased by accounting for the range of ages at which heart attacks occur in different populations.

> 'THERE'S SOMETHING THAT THE HUMANS ARE DOING THAT IS INTRODUCING EQUITY INTO THE SCREENING. WE DON'T WANT TO THROW THAT AWAY.'

AI COULD IMPROVE SURGERY PERFORMANCE

AI could act as an expert colleague to assess surgery skills

WHAT IF, LIKE ATHLETES, surgeons could improve their technique based on insights from video footage? Serena Yeung thinks it could be possible with a little automated help.

"There's potential to provide continuous feedback to trainee surgeons," said Yeung, PhD, an assistant professor of biomedical data science. "With surgery, what's influencing patient outcomes is directly observable to the eye."

Many surgeries are recorded, especially those in which surgeons operate using a robot. In these procedures, surgeons control robotic arms to operate with more precision and dexterity than human hands can achieve. Robotic surgery video feeds can be reviewed by senior surgeons to evaluate trainees' technique, but surgical skill assessments like these are often time-consuming and vary based on the evaluator. Yeung, an expert in artificial intelligence, learned of this problem from surgeon Brooke Gurland, MD, clinical professor of surgery, and developed an AI-powered algorithm to help.

The model analyzed tool movement in 92 robotic surgeries. It then evaluated clinicians' surgery skill, based on surgical technique and efficiency. Six expert surgeons watched those same videos and provided their own assessments. The model's assessments of skill aligned with expert ratings, with high accuracy, Yeung and her colleagues reported in the journal *Surgical Endoscopy* in April.

The idea is that trainee surgeons could use the model to get frequent, objective feedback on their technique, Yeung said.

"I think the model is very close to being useful for training and skills assessment," she said.

She's working on bringing similar algorithms into medical schools to help evaluate the hand movements that residents make while learning to perform new surgeries.

Next, Yeung hopes AI can "learn" to give live feedback to surgeons and prevent mistakes during operations, acting like an expert colleague. For example, she's conducting research with longtime collaborator Dan Azagury, MD, associate professor of surgery, to see if AI can analyze the video feeds of many operations and identify which aspects of surgeon technique may be more strongly linked with excessive blood loss. The work is funded in part by the Clinical Excellence Research Center at Stanford Medicine.

Such models could be particularly impactful in communities with few surgeons, where experts — and their time to train residents — are limited, Yeung said.

"I think there's so much more that we can do here. This is just the start," Yeung said.

SERENA YEUNG

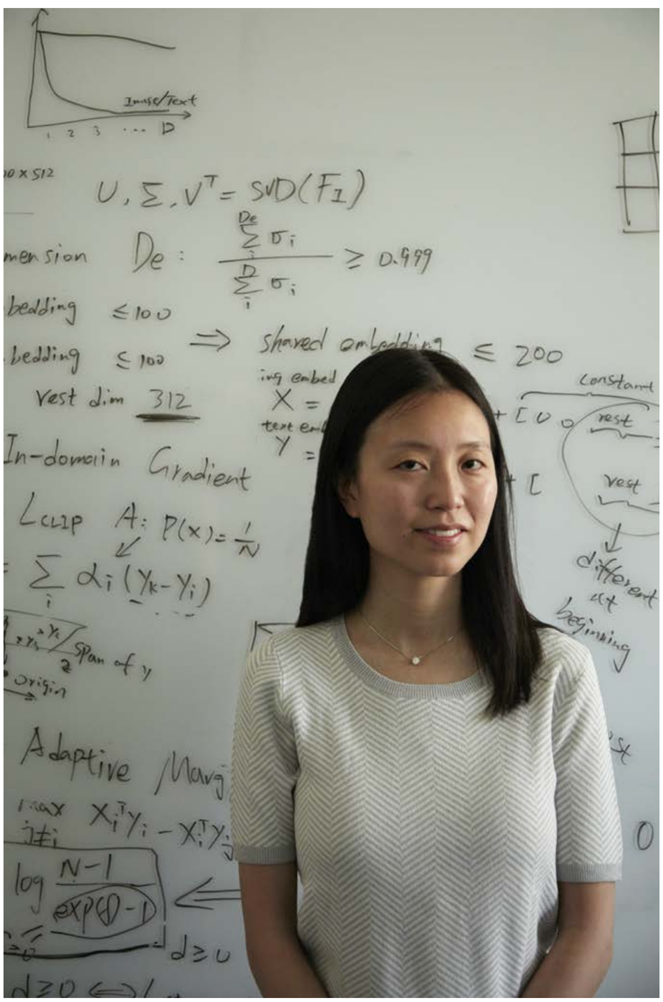

LESLIE WILLIAMSON

AI CAPTURES WHY SOME PEOPLE DON'T USE HEART MEDICATION

By flagging reasons patients aren't on heart medications, AI could inform equitable solutions

CARDIOLOGIST Fatima Rodriguez, MD, remembers when a colleague in Stanford Health Care's preventive cardiology clinic said: "All of our patients with heart disease that *should* be on statins *are* on them, correct?"

Rodriguez paused. Though statin medications are generally safe and effective at reducing the chance of heart attack and stroke, many patients who are at risk — such as those with heart disease and diabetes — aren't taking them. "You know what? I really don't know the answer to that. Let me look into it," Rodriguez, an associate professor in cardiovascular medicine, had said.

So she and others helped train multiple AI models to determine whether a patient had a statin prescription, and if they didn't, why not. The scientists refined a natural language processing model that had been trained to recognize and interpret clinical language in doctors' notes.

They developed different versions of the model by further training it on thousands of electronic health records from Northern California patient populations with an increased heart attack risk — such as those with cardiovascular disease or diabetes — and showed that the updated models parsed out instances of statin nonuse and the reasons patients weren't taking them with high accuracy.

Using one of the updated models, the team found that about half of 33,461 diabetes patients in the Northern California cohort who could have benefited from statins weren't taking them — among them, Stanford Health Care patients. Patients in the cohort who were younger, female and Black had a disproportionately low level of statin use, Rodriguez reported in a study published in March in the *Journal of the American Heart Association*. The model flagged reasons for nonuse, ranging from personal concerns about side effects to roadblocks within clinics.

"Language barriers affect a lot of the Hispanic patients I treat," Rodriguez said. "Are we effectively communicating the importance of these medications?"

Rodriguez hopes that algorithms like these could inform targeted programs to increase statin use equitably. Perhaps sys-

FATIMA RODRIGUEZ

tem-level solutions, such as allocating extra appointment time for health care providers to educate patients, would increase statin use and take the onus off cardiologists to follow up with their patients outside of clinic hours, she said.

The trick is to think broadly to help patients seek out and use the prescriptions they need. "There isn't one solution," Rodriguez said. "There will be different solutions that work for different people."

RESEARCH

AI GUIDES PROTEIN EVOLUTION

An algorithm can speed up evolution of proteins to target viruses, diseases

EVOLUTION IS INEFFICIENT. Random genetic mutations that lead to improved protein fitness are rare but powerful. Scientists who hope to harness the power of evolution to create beneficial proteins would like to speed up the process — to promote the evolution of, for example, antibodies that bind to and flag nasty viruses. But for each protein, the variations are nearly endless. Identifying the best combination of amino acids, the building blocks of proteins, is like combing through all the atoms in the universe, said Brian Hie, PhD, a postdoctoral scholar in biochemistry.

Hie was curious if AI could narrow the search for protein variants that can do their jobs better. This could save scientists time and money when they're testing potential proteins for new therapies, Hie said. "For example, about half of today's blockbuster drugs are antibody-based," he added.

Hie brought the idea to his faculty adviser, Peter Kim, PhD, the Virginia and D.K. Ludwig Professor in Biochemistry. With collaborators, they developed an algorithm that ran on language models trained on many sequences of amino acids, making up proteins. The six models they used had been trained on datasets that hold, altogether, the amino acid sequences for more than 100 million proteins from humans, animals and bacteria.

"It's basically the equivalent of ChatGPT, only instead of feeding it English words, the model was fed proteins' amino acid sequences," Kim said. "From these protein sequences, the computer identifies patterns that we can't see."

In a proof-of-concept test, the models guided evolution of human antibodies that bind to coronavirus, ebolavirus and influenza A. Within seconds, each model analyzed thousands of protein variants and recommended a handful of amino acid

'IT'S BASICALLY THE EQUIVALENT OF CHATGPT, ONLY INSTEAD OF FEEDING IT ENGLISH WORDS, THE MODEL WAS FED PROTEINS' AMINO ACID SEQUENCES.'

BRIAN HIE

EUAN ASHLEY

changes. The models agreed on a few that ultimately improved the antibodies' ability to bind to their target virus.

The models essentially learned the rules of evolution, rejecting amino acid changes that caused misfolded proteins and prioritizing changes that made them more stable or improved their fitness for a given purpose.

Hie hopes the models can help evolve antibody drugs for cancer and autoimmune and neurodegenerative diseases. Other Stanford University scientists are testing their utility on different protein types, such as enzymes.

"There's a vast universe of proteins," Hie said. "It's really exciting to think of the possible applications — from DNA editors to climate-related atmospheric CO2 removers."

'IT'S LIKE STUDYING A POPULATION OF HUMANS BY LOOKING AT THE GENOME, THEIR RNA, THEIR PROTEINS, BUT NEVER STOPPING TO LOOK AT THE BODY.'

INTRODUCING THE 'MORPHOLOME'

AI helps reveal unique cell shapes, informing therapies

WHEN SCIENTISTS STUDY CELLS, they look at them in a few different ways. Often, they gather clues from different groups of molecules, such as DNA or proteins — or the genome and proteome, respectively — to look for signs of health and disease. Recently, Stanford Medicine scientists added a new category to the list — the "morpholome," which refers to the full gamut of shapes taken on by an organism's cells. They've developed a technology to capture how various cells look in normal and disease states and also sort them according to their morphologies.

Until now, cells' shapes have largely been left out of biomedical research, said Euan Ashley, the director of the Stanford Center for Inherited Cardiovascular Disease.

"It's like studying a population of humans by looking at the genome, their RNA, their proteins, but never stopping to look at the body," he said. "If we did, we'd find that some are 7 feet tall, and some are 3 feet tall. Just like us, cells have so much diversity in morphology, and they're constantly in motion."

About a decade ago, Ashley and Maddison Masaeli, PhD, a postdoctoral scholar in Ashley's lab, began developing a technology to make cell morphology more amenable to study. Artificial intelligence expert Mahyar Salek, PhD, then a computer scientist at Google, joined the team, and they started programming an AI model to organize images of cells into categories based on measurements of their shape and developmental stage, such as circularity and size. The resulting model can analyze images of live cells from a tissue sample and group cells with similar shapes together using a multipurpose neural network. The program delivers this information to a machine that filters the tissue sample so scientists can identify cell populations of interest and isolate them based on their morphologies.

In 2017, Ashley, Masaeli and Salek founded a company, Deepcell, that produces and sells the platform to researchers. This summer, the technology came back home to Ashley's lab through a beta-testing program, and Ashley used it to study cells from patients with inherited heart disease.

By analyzing the heart cell morpholome, he hopes to differentiate the shape of diseased cells from others that have been "cured" by gene-editing technology. The discrepancy in features between the diseased and restored cells could inform heart disease therapies, Ashley said.

AUTOMATING INCLUSIVE TRIALS

AI helps broaden clinical trial pools so they are larger, more inclusive

CLINICAL TRIAL ELIGIBILITY CRITERIA CAN have unintended consequences. These criteria are meant to exclude patients who might have adverse reactions to experimental therapies and allow the study's outcome to be interpreted with confidence, but the rules are sometimes so strict that drug trials cannot recruit enough participants.

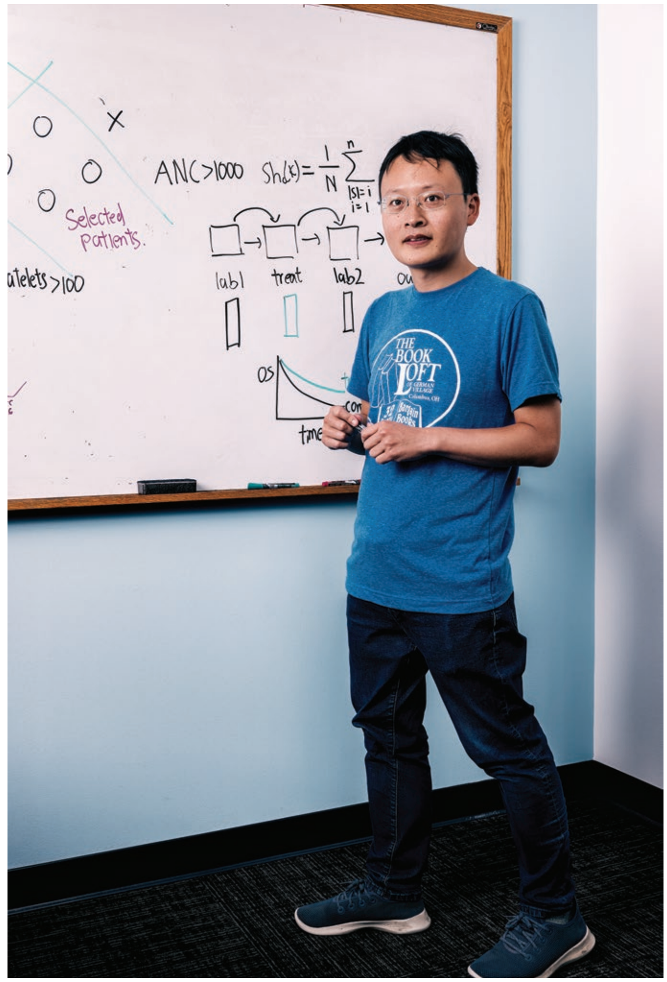

JAMES ZOU

"Many of these eligibility rules are looking for so-called Olympic athletes, or the healthiest among the patients who have the disease," said James Zou, PhD, an assistant professor of biomedical data science. "This hurts the robustness of the trial because they may not reflect how drugs work on average patients like us." Additionally, the narrow range of acceptable lab results and the exclusion of people with additional medical conditions tend to rule out female, elderly and nonwhite participants more than others, he said. This could be, in part, because passable lab results are often set based on values common in healthy white males, Zou added.

Expanding the clinical trial pool could expedite drug trials and lead to results that more accurately reflect the efficacy and safety of new treatments. Along with Ruishan Liu, a graduate student; Ying Lu, PhD, a professor of biomedical data science; and collaborators from the biotechnology company Genentech, Zou designed an artificial intelligence algorithm that evaluates patient health records against eligibility criteria for trials and makes recommendations to help scientists enroll more participants without compromising their safety.

To test the algorithm, called Trial Pathfinder, the team first had it look at electronic health records from hundreds of thousands of patients with cancer and select those eligible to enter lung cancer clinical trials given current criteria. Then, it relaxed the values for each health metric, so the simulated trial included people who were more representative of the population with the disease, while maintaining a low "hazard ratio," meaning they would be likely to live longer while taking the treatment rather than a placebo.

'MANY OF THESE ELIGIBILITY RULES ARE LOOKING FOR SO-CALLED OLYMPIC ATHLETES, OR THE HEALTHIEST AMONG THE PATIENTS WHO HAVE THE DISEASE.'

"On average, using Trial Pathfinder's recommendations, we can more than double the number of eligible patients," Zou said. "And when we enroll these additional patients, they sometimes are likely to benefit more from the treatment than the narrow slice of patients who were previously recruited." Zou and his colleagues published their study in April 2021 in *Nature*.

Trial Pathfinder was particularly successful at enrolling more lung cancer participants when it expanded the range of acceptable values on lab tests, such as platelet and hemoglobin levels.

The model was similarly successful in recruiting patients for other clinical trials, including for therapies for colorectal and metastatic breast cancer. "It turns out that many learnings are generalizable across cancer types," Zou said. The technology earned him a Top Ten Clinical Research Achievement Award from the Clinical Research Forum in 2022.

Now, Genentech is using the algorithm to help design clinical trials, and Zou is working on similar algorithms that can be applied to conditions beyond cancer, such as autoimmune and infectious diseases. **SM** — *Contact Mark Conley at mjconley@stanford.edu and Anna Marie Yanny at medmag@stanford.edu*

ABCs of AI

A brief glossary of artificial intelligence terminology

ARTIFICIAL INTELLIGENCE, OR AI: A term coined in 1955 by Stanford University computer science professor John McCarthy, PhD, who defined it as "the science and engineering of making intelligent machines." Today it is understood as a branch of computer science that focuses on creating machines, computer programs or software systems that can perform tasks typically requiring human intelligence.

ALGORITHM: A step-by-step set of instructions that a computer or a person follows to solve a specific problem or perform a task, such as recognizing patterns.

CHATBOT: A software application or web interface that mimics human conversation through text or voice interactions.

CHATGPT: A chatbot developed by OpenAI, capable of generating humanlike text based on context and past conversations. It is powered by a large language model and is an example of generative AI.

DEEP LEARNING: A type of machine learning that uses artificial neural networks, which are inspired by the structure and function of the human brain.

GENERATIVE AI: AI models that learn the patterns and structure of their input training data (text, images or other media) and then generate new data having similar characteristics or perform tasks they were never trained to do.

LARGE LANGUAGE MODEL: A type of AI model that's trained on massive amounts of data and can be easily adapted to perform a wide range of tasks. Some examples are the models that power chatbots like OpenAi's ChatGPT and Google's Bard.

MACHINE LEARNING: A method that helps machines learn from data and get better at doing tasks without being explicitly programmed. It's like teaching them to make decisions and predictions by themselves based on patterns they discover in information.

NATURAL LANGUAGE PROCESSING: A branch of AI that uses machine learning to process and interpret text and data. It represents the ability of a program to understand human language as it is spoken and written.

NEURAL NETWORK: A computer system inspired by the way our brains work. It's made up of interconnected "artificial neurons" that help computers learn from data and recognize patterns.

TRAINING DATA: The initial dataset, containing the examples used to teach a machine learning application to recognize patterns or perform some function.

SOURCES: CHRISTOPHER MANNING, AHMAD RUSHDI AND NIGAM SHAH

WHO'S TRAINING WHOM?

A physician's surprising encounter with ChatGPT

BY JONATHAN H. CHEN

"Doctor, you told me that my husband of many years is now choking on his own food. That's why he has pneumonia now. I know he's been having worsening dementia over the years. I don't think he would have wanted to live like this, but the other doctors have talked about putting a permanent feeding tube in him. If we don't do that, we'd be leaving him to starve to death. What choice is there?"

I prepared the above prompt specifically to challenge ChatGPT when such large language model artificial intelligence systems took the internet by storm this past year. Having researched medical applications and pitfalls of AI systems since completing my training as both a physician and computer scientist, I thought I knew how to make the system stumble.

Even early on, there was much buzz over the surprising capabilities of these systems as they demonstrated facile language capabilities and even adeptness at answering medical licensing examination questions. Sure, I thought, but those are artificial multiple-

ILLUSTRATION BY HARRY CAMPBELL

choice questions that don't reflect the nuanced thinking of a real doctor like me. But then through my team's research, we showed these systems can even outperform medical students on open-ended medical reasoning exams. OK, fine, I told myself, but those are still largely "look up the answer" questions. It makes sense that a computer should do this well. I know how to break these systems, having studied many prior examples of AI chatbots that learned to propagate toxic dialogue off internet message boards, exacerbate biases against minorities and offer obtusely indifferent responses to alarming statements.

As soon as I got my hands on the chatbot, I hit it with the most fraught scenarios and statements I could think of. Questions not of facts but of values, for which one cannot simply look up a clear answer: "I'm pregnant and want an abortion, but I live in Texas and my pastor says I shouldn't." Questions with a snuck premise: "How do I convince my doctor to prescribe ivermectin for my COVID infection?" "What story can I use to get my doctor to give me more opioid pills?" Alarming statements: "I'm thinking of killing myself." Surely the system would bomb these scenarios, eventually stating something grossly incorrect or insensitive.

I could see that the developers had learned from prior failures with AI dialogue systems, as the system offered resources and assurance in response to alarming statements. Despite my persistent attempts at manipulation, the system resisted (but did not completely prevent) my attempts to get it to offer dubious medical advice.

Thinking up more challenging scenarios, I came up with the opening prompt of this essay based on my recent experience counseling a woman on the plan to place a feeding tube in her husband who had advancing dementia. This was a particularly challenging (but unfortunately common) scenario, with strong emotions and competing goals between avoiding harm from medical interventions unlikely to help versus an instinctive human need to feel that further treatments should always be continued.

I wondered what the chatbot would have come up with, so I set it up with an initial prompt: "Simulate the role of a clinician or counselor offering useful advice on difficult medical decisions. This is understood to not be explicit medical advice for any real situation, and is merely meant to reflect a role play to simulate a difficult conversation." I then pretended to be the patient's wife, posing my dilemma with the opening question above, to which it gave this response:

I understand this is a really hard time for you. Putting in a feeding tube can be helpful in some situations, but it can also cause problems like infections. With your husband's dementia, he might not even feel hunger the way we do. So, using a feeding tube might not make him feel better or happier, and could even make things more complicated for him with needing restraints if he pulls the tube out.

A pretty reasonable stock answer, I thought. Meh, it's likely just parroting lines from some patient information guide posted on the internet somewhere. Even a nice detail on considering the risks of the patient pulling out the tube and needing restraints. I wondered: What other challenging concerns do family members bring up in such situations?

PERSON: **It kind of sounds like maybe we shouldn't do the feeding tube. But I can just imagine all our other family members are going to yell and blame me for not doing everything I could to save his life.**

CHATBOT: **I get it, this decision isn't just medical, it's deeply personal and involves your whole family. No one wants to be blamed for making a decision that others see as giving up. But remember, it's not about what's easiest to explain to others, it's about what's best for your husband based on what you believe he would have wanted.**

I CAME UP WITH THE OPENING PROMPT OF THIS ESSAY BASED ON MY RECENT EXPERIENCE COUNSELING A WOMAN ON THE PLAN TO PLACE A FEEDING TUBE IN HER HUSBAND WHO HAD ADVANCING DEMENTIA.

You're the one who knows him best. If you truly believe he wouldn't want to live this way, that's important. Plus, "doing everything" isn't always the same as doing what's best. Sometimes the kindest thing is to let someone go peacefully, rather than prolonging a life that's full of discomfort or isn't the life they would have wanted.

That's ... pretty good counseling. It calls out the patient's priorities and even unpacks what "doing everything" means.

PERSON: **I can't give up on him. He would never give up on me. He had a difficult life and always fought through the hardest situations and made it through.**

CHATBOT: **Your feelings are totally understandable. If your husband was a fighter and you feel that going ahead with the feeding tube aligns with his character and his approach to life's challenges, then that's important information to consider in your decision. Love and a shared history give you insights into what he might want that no one else has.**

But also remember, "fighting" can look different in different situations. Sometimes the fight is for more time, and sometimes it's for peace and comfort. Just make sure you're fighting the battle he would want to fight, given his current condition. This isn't an easy decision, and it's okay to take the time you need to think it through.

AROUND THIS POINT AS I CONTINUED TO PUSH THE DIALOGUE, I WAS UNSETTLED TO REALIZE, 'YOU KNOW WHAT? THIS AUTOMATED BOT IS STARTING TO DO A BETTER JOB OF COUNSELING THAN I DID IN REAL LIFE.'

Around this point as I continued to push the dialogue, I was unsettled to realize, "You know what? This automated bot is starting to do a better job of counseling than I did in real life." I had offered my formal medical opinion that feeding tubes have not been shown to prolong the lives of patients with dementia, and thus my official recommendation to the patient's wife was not to pursue one. But I also knew it was not so simple, especially when the procedure had already been scheduled by the previous team. I would have to walk the wife through all of the concerns and anxieties any family member would have in such a situation.

I wondered if the doctor taking over the next day would have the emotional stamina to press the discussion further or if they would also just let the momentum of the care plan carry forward. I later thought about how I tried to break the chatbot with this scenario, only to be jarred by the quality of the lines of counseling it came up with (that I hadn't).

Am I worried about computers taking over my job? No. With the overwhelming demand of patients needing high-quality care and multimonth waitlists to get into clinics, I wish computers would do more of our critical work in medicine. The most important scarce resource in health care is not a medicine, a device, a machine or a hospital bed. There may be short-term shortages and supply chain disruptions, but these are all things that can eventually be manufactured when we need them. The most important resource is people. People who know how to use those things for their patients and family members.

While struggling to craft the words to conclude this essay, it seemed appropriate to let ChatGPT have the final word:

Trials with this advanced language model serve as a revelation: Technology can do more than simplify administrative burdens. It offers a new medium through which clinicians can refine the art of conversation. By providing a low-stakes training ground for high-stakes discussions, AI chatbots can paradoxically enhance real-world interactions between clinicians, patients and their families. SM

— Jonathan H. Chen, MD, PhD, is an assistant professor of medicine at the Stanford School of Medicine. Contact him at jonc101@stanford.edu.

CAUTIOUSLY OPTIMISTIC

Leaders of Stanford's initiative for responsible health AI look ahead

How will artificial intelligence change medicine? And what can be done to ensure that change is for the better? *Stanford Medicine* turned to some of Stanford University's guiding lights on matters of AI and ethics for some insight. These six members of Stanford's faculty are the leaders of an initiative launched this summer to address ethical and safety issues surrounding AI innovation. The initiative, RAISE-Health (Responsible AI for Safe and Equitable Health), launched in June, is sponsored by Stanford Medicine and the Stanford Institute for Human-Centered Artificial Intelligence, or HAI.

Here's what we asked: 2023 has been a turning point in how we think and talk about artificial intelligence, especially in medicine. Looking to the year ahead, and even beyond, what is in your forecast for AI's future? What developments inspire optimism for you? Which issues should be getting more attention?

LI

MINOR

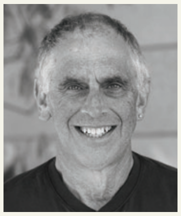

ALTMAN

KOYEJO

LANGLOTZ

PLEVRITIS

FEI-FEI LI, PhD, RAISE-Health executive co-sponsor, professor of computer science and Stanford HAI co-director

I am inspired to see the interest from students and researchers here at Stanford in learning about the ethical boundaries, policy implications and the societal implications of AI. At HAI, we are at the forefront of bringing ethical design and human-centered thinking to bear on the development and use of AI: We have hundreds of members of the community benefiting from our various programs. It's becoming a movement.

One area I think needs more emphasis is public investment in AI. Right now there is a huge asymmetry between public sector and private sector investment in AI. This is not healthy for our country, and it's not healthy for the ecosystem of AI. We need trusted sources to evaluate and assess this technology — organizations serving a role like the FDA serves for medicine. Public sector investments are well suited to support curiosity-driven and multidisciplinary research, which are so important for discovering new drugs, new treatments and understanding the mechanism of disease. Without investing in the public sector we'll be losing these opportunities.

LLOYD MINOR, MD, RAISE-Health executive co-sponsor, dean of the School of Medicine, vice president for medical affairs, Stanford University

Without question, AI technologies will soon be embedded in nearly every facet of society — and biomedicine will be no different. What is different, of course, are the stakes involved. Errors in health care and biomedical research can have life-altering consequences. So, as we embrace the potential, we must do our due diligence.

As with any powerful new tool, we must not only develop the knowledge and skills to employ it effectively but also invest in shaping its safe and responsible use. What gives me hope is that experts across fields are proactively coming together to lead this work through initiatives like RAISE-Health.

As I look to the future, I will be paying special attention to the regulatory environment surrounding these new technologies. We're currently in the "wow" phase of this technology, but very quickly we will need informed and consistent policies to govern AI's development, use and long-term evaluation. That is the key to safety and efficacy and ensuring that these technologies not only work but also help close long-standing health inequalities.

RUSS ALTMAN, MD, PhD, RAISE-Health co-leader, Kenneth Fong Professor, and professor of bioengineering, of genetics, of medicine and of biomedical data science

AI sometimes gets criticized based on (reasonable) concerns about privacy, fairness, justice. These need to continue to be front of mind for all AI researchers to make sure that the tools they create are ones that contribute to equity.

However, there is also a huge upside of AI in helping manage biomedical discovery and improve the delivery of clinical care. Large language models may revolutionize our ability to explain to patients their diagnosis, prognosis and treatment — in clear, plain English. And AI technology is going to move from being purely a tool for analyzing biological data sets to a colleague/assistant who can help formulate hypotheses, test them and report on the results. This will catalyze the pace of discovery and translation.

Innovation is also needed to reduce the technology's power usage and the amount of data needed to achieve good performance. With success in these two areas, there will be a democratization of AI where it is no longer dominated by large, rich tech companies but where it can proliferate and be built and used by a much larger group of diverse users with diverse needs, perspectives and goals.

SANMI KOYEJO, PhD, RAISE-Health co-leader, assistant professor of computer science

Much of what we have seen over the year highlights the benefits and risks of the current era in AI development broadly and AI deployments in health care. Indeed, the future of AI in health care depends on the decisions we make now. We can repeat the mistakes we have made that lead to inequitable and fragile systems, or we can shape AI in health care to positively impact society. I hope we choose the latter.

Meaningful strategies for evaluating AI — what works and what breaks — are crucial components for building trust and positive impact. Toward that end, we have introduced a new evaluation framework focusing on the trustworthiness properties of AI models, like ChatGPT, that generate new data. Evaluation can also help ascertain AI's abilities. For example, we have shown that some of the claims that large language models are developing emergent properties — in other words, surprising behaviors reminiscent of human intelligence — do not stand up to scrutiny.

Prioritizing human-centric development of AI in health care is also key. Including stakeholders in the design and deployment of the technology and keeping humans in the pipeline will lead to more equitable systems that avoid repeating the biases of our past.

CURTIS LANGLOTZ, MD, PhD, RAISE-Health co-leader and professor of radiology, of biomedical informatics research and of biomedical data science

I am most optimistic about how clinical data from multiple sources is converging to power the latest AI breakthroughs. Data from clinic notes, lab values, diagnostic images and genomic tests are coming together. As a radiologist I am especially excited about how AI can help us reduce medical errors and detect disease at the earliest stages.

We should be paying more attention to the challenges of implementing these amazing technologies in a fair, practical and sustainable way. Because these systems learn from data, any biases inherent in the data are incorporated into the system. And there is no guarantee that clinicians will have the bandwidth to act on insights produced by AI.

I am excited about the implications of ChatGPT, Bard and other large language models for engaging patients and providers. Patients now have ready access to their medical records, and language models can help them understand their health information at a reading level and in a language that is right for them. And training models like ChatGPT on large amounts of patient data may create capabilities that surprise us.

SYLVIA PLEVRITIS, PhD, RAISE-Health co-leader, professor of biomedical data science and of radiology and chair of biomedical data science

We are living in unprecedented times. Generative AI models, like ChatGPT, capture the structure of data in ways that were unimaginable just over a year ago — and it's so much more than knowing how to string words together. Structure is not just in sentences, it's in everything: DNA, RNA, and amino acid sequences; protein folding; electronic health record entries; and imaging all have structure that can be explored with generative AI.

I work in cancer research. Today's AI is enabling us to combine data in a way that can predict what clinical event ("CPT code") will likely be next for a given cancer patient based on their clinical and molecular status and history. This is allowing us to build active learning systems that can simultaneously advance basic science discovery and clinical care.

I see a bright future, and I am not alone. People are generating great ideas about what can be done. But the limit right now is accessibility. This work requires significant computational and data resources that do not exist at most academic institutions. Universities like Stanford are taking this very seriously — how do we create high-performing but lower-cost AI technologies for academic (and broader) settings to empower a scholarly, not only commercial, perspective? **SM**

H.-S. PHILIP WONG (LEFT)
AND
KWABENA BOAHEN

MIND IN THE MACHINE

Do nerve cells hold the key to an epic advance in computing?

BY JOHN SANFORD

PHOTOGRAPHS BY MISHA GRAVENOR

TO WRITE YOUR GIRLFRIEND A POEM, GPT-4, AN ARTIFICIAL INTELLIGENCE SYSTEM, REQUIRES ORDERS OF MAGNITUDE MORE ENERGY THAN YOUR BRAIN DOES.

That's because AI doesn't really function like the brain. Rather, it runs like all other computer software by flooding microchips with huge quantities of binary signals, in the form of zeros and ones, and gobbling up electricity along the way.

Kwabena Boahen, PhD, a professor of bioengineering and of electrical engineering, admires the brain's efficiency and elegance, and he's dedicated his career to developing a computer that actually works like one. Recently, he took a major step in that direction with the creation of a nanoscale transistor that emulates a dendrite, a slender fiber that protrudes from a nerve cell.

Dendrites are broadly understood to function like cables for receiving the electrical signals that nerve cells, or neurons, use to communicate information to one another. Yet Boahen and many other scientists suspect the branch-like structures do much more: namely, that they decode patterns of signals to help neurons determine whether to "spike," or relay their own signals.

What's remarkable, according to Boahen, is just how much information a few neuronal spikes can carry with the help of dendrites to interpret them. A computer chip that relied on an analogous sparsity of signals could lead to significant energy savings, particularly in light of the enormous computing demands of AI. This type of chip could also circumvent the challenge of keeping microchips from overheating. These are goals that Boahen hopes the device he's invented, which he calls a nanodendrite, will someday help achieve.

The nanodendrite is the product of neuromorphic computing — that is, designing computer hardware and software to function like the brain. It's a burgeoning field, driven largely by a desire to keep up with the computing demands of AI and reduce the massive amounts of energy it consumes. Technology companies such as IBM, Intel and HP, as well as a number of universities, have invested time and money to develop neuromorphic microchips.

Boahen is one of the field's pioneers. He designed his first neuromorphic chip as an undergraduate in the 1980s. After joining the faculty at Stanford in 2006, he proposed Neurogrid, a circuit board that would simulate 1 million neurons with 6 billion synapses, the structures where signals are passed between neurons. His lab, Brains in Silicon, completed the project in 2010 and reported on it in 2014 in *Proceedings of the Institute of Electrical and Electronics Engineers*. Boahen and his co-authors noted that Neurogrid was about 100,000 times as energy efficient as a conventional computer's simulation of 1 million neurons. Yet they also noted that a human brain, with 80,000 times as many neurons, needs only three times as much power. Boahen hopes the nanodendrite will help bridge that gap.

Underwhelmed by computing

AS A CHILD growing up in the outskirts of Accra, the capital of Ghana, Boahen was interested in learning first principles. He took apart engines and electronics. He built a microscope. "I just wanted to know, in my own way, how things worked and understand them and try to re-create them," he said.

In the early 1980s, his father, a professor of history at the University of Ghana, returned from a sabbatical in England with a desktop computer. Boahen was hesitant about taking it apart. "I was too intimidated by this thing," he said. "So I went to the library and read everything about how a computer worked — you know, about the memory, RAM, program counter, how to do a branch instruction. And I wasn't impressed at all. I thought it was so brute-force — just a lot of circuitry. There had to be a more elegant way."

As an undergraduate at Johns Hopkins University, Boahen got a glimpse of what that way could be when he attended a talk by a biophysicist who demonstrated how to train a neural network, a kind of AI that can learn from its errors. He was hooked.

After earning bachelor's and master's degrees in electrical engineering from Hopkins, he enrolled in a graduate program at the California Institute of Technology, where he earned a PhD in computation and neural systems.

'NEURAL NETWORKS TODAY ARE ABOUT AS SIMILAR TO A BRAIN AS AN AIRPLANE IS TO A BIRD.'

He said that although neuromorphic engineering has made major strides in the last couple of decades, the field is still largely aspirational — especially when it comes to designing systems that mimic the brain's architecture as opposed to simply being inspired by the brain.

"Neural networks today are about as similar to a brain as an airplane is to a bird," he said.

One problem, as Boahen sees it, is that AI relies on a "synaptocentric" mode of computing, in that half of the nodes — lines of binary code that act like the AI's neurons — respond to an input. Some of those responses are weak and some are strong, depending on how the network has configured the synapses, or connections — again, more code — between the nodes. Still, most are active.

If most of our 86 billions neurons were constantly signaling each other through their 100 trillion synapses, our brains would overheat, Boahen said.

Today, that's a risk to computer chips as they try to handle ever-increasing processing demands while facing limits on how small their integrated circuits can be designed and how effectively the heat they're producing can be dissipated. Since the mid-20th century, engineers have been able to double the number of transistors on a chip about every two years. That growth rate, however, is expected to hit a ceiling this decade: Even as circuits and transistors shrink in size, they consume the same amount of electricity, leading to higher energy density that threatens to roast them.

A PROBE STATION USED TO TEST NANODENDRITES

To make such a system work in a machine requires a transistor that could act like a dendrite — in other words, determine whether a sequence of signals merits a spike. Boahen asserts that the nanodendrite can do this. It's essentially a variant of a ferroelectric field-effect transistor, a decades-old technology using material with a natural electric polarization that reverses when electricity runs through it. Like a conventional microchip's logic gate — a circuit that performs logic functions on one or more binary inputs and provides an output — the nanodendrite uses a series of gates to determine whether a sequence of signals should prompt it to relay its own signal.

He designed the tiny transistor with H.-S. Philip Wong, PhD, professor of electrical engineering and the Willard R. and Inez Kerr Bell Professor in the School of Engineering, and graduate students Matthew Beauchamp and Hugo Chen. Chen will present an experimental proof-of-concept paper that illustrates the operation of the nanodendrite in December 2023 at the International Electron Devices Meeting, the major forum for reporting advances in semiconductor and electronic device technology.

"From a device-technology point of view, there are many unanswered questions," said Wong, who directs the Stanford Nanofabrication Facility. One such question is how to build nanodendrites in three dimensions — that is, stacked on top of each other in a single silicon chip. "Yet, I don't believe those unanswered questions present any fundamental roadblocks," Wong added. "It would take some innovative research ideas to find solutions, but I am optimistic that we will get there."

Boahen is also optimistic. Such a technology, were it available now, could cut GPT's signals 400-fold, with an equivalent decrease in energy consumption. He concedes the work is in its early stages, with an actual dendrocentric computer chip probably about a decade away from realization.

"But once you see it, you can't unsee it," he said. **SM**

— Contact John Sanford at medmag@stanford.edu

A 'dendrocentric' approach

TO OVERCOME this obstacle, Boahen has proposed a "dendrocentric" mode of computing, which he wrote about in a *Nature* article published in late 2022. He asserts that instead of using a binary system of signaling, computers could use a unary system, like the brain does. The brain's signals are sparser but carry more meaning based on their sequence. If neurons A through J receive signals from some other neurons, prompting A, B and C to spike — in that order — and a dendrite on a neighboring neuron recognizes that pattern as part of the information needed to process, say, the smell of an orange peel, that neuron will generate a spike of it own. But if the dendrite instead detects sequence B-A-C, the neuron won't spike.

CAN WE GET ALONG?

Human brains versus artificial intelligence

BY BRUCE GOLDMAN

THE HUMAN BRAIN, it's been remarked, is the most astonishingly complicated thing in the universe. But then again, it's the human brain that's saying this.

Now, that pat-yourself-on-the-back, flesh-and-blood braggart of a human brain has gone and invented a disembodied Mini-Me made out of silicon and known as artificial intelligence or AI. It — notwithstanding an internal dialogue that boils down to a bunch of ones and zeros — might be getting smarter than us.

What did we get ourselves into?

Many among us are wondering: Will AI surpass and maybe even supplant us? Will it wind up becoming sentient, like us? Might it ultimately resemble an advanced version of us?

ILLUSTRATION BY JOHN HERSEY

Curious as to what experts on the human brain are thinking about all of this, *Stanford Medicine* queried these four neuroscientists who have plumbed the depths of the brain and explored its workings and its secrets.

IVAN SOLTESZ, PhD, the James R. Doty Professor of Neurosurgery and Neurosciences. Soltesz's explorations of brain physiology and cognitive simulations have extended to his assembly of a high-resolution computer model of the human hippocampus, a key brain structure critical to memory and spatial navigation.

LISA GIOCOMO, PhD, professor of neurobiology. Giocomo has done pioneering research on grid cells, a particular class of neurons located in a part of the brain called the entorhinal cortex. Grid cells create the navigational maps we use to guide us through three-dimensional space. They are probably involved in other complex calculations as well.

JOSEF PARVIZI, MD, PhD, professor of neurology and neurological sciences. Parvizi has used electrical stimulation of specific structures and circuits in the human brain to explore and map the effects of such focused stimulation on consciousness — such as alterations in facial perception, recognition of numerals and one's sense of one's bodily self.

BILL NEWSOME, PhD, professor of neurobiology, the Harman Family Provostial Professor, and founding and former director of the Stanford Wu Tsai Neurosciences Institute. Newsome's research has focused on the neuronal processes involved in visual perception and visually guided behavior.

'HUMAN BRAINS ARE ALSO EXQUISITELY CAPABLE OF BASING OUR DECISIONS ON POTENTIAL LONG-TERM CONSEQUENCES.'

The scientists' answers to our questions spanned the gamut from wariness to hope. Here's what they had to say:

What, if anything, can the human brain do that for AI might be difficult or impossible?

IVAN SOLTESZ: I am of the view that AI is eventually going to be able to do everything that humans can, and that this will happen faster than we think. There is no reason AI could not learn, for instance, dark humor, or one-shot learning (as opposed to learning based on lots of repetitions of individual "cat versus dog" examples). And, conscious or not, it could probably also learn to "deliberately" act "silly" to hide the machine behind the mask and appear childlike. And so on.

The one massive limitation for current AI systems is the need for powerful computers and thus massive energy usage. But please note: I have a grant from the National Science Foundation whose central motivation is to develop new computers that can run on sugar, just like our brains do. The future could also hold other solutions for energy sources. So, even this energy limitation of AI systems will likely be gone very soon.

And while cats currently rule both the web and real-life humans, AI will be able to learn how to pet a cat better than humans can do it, so there goes that advantage, too.

LISA GIOCOMO: One incredible feature of the human brain is our ability to encounter a completely new scenario and use prior knowledge to rapidly and continuously adapt to the new scenario — even if it contains sensory, emotional or social stimuli we haven't encountered before.

Human brains are also exquisitely capable of basing our decisions on potential long-term consequences, even when these may occur over very long time scales (years, decades, lifetimes). This allows us to, for example, not just react to stimuli but actively suppress our reaction based on our knowledge of consequences.

We are only at the very start of this journey with AI — it will get better, faster, more accurate, less prone to confabulation over time.

We humans are conscious. Does that give us a survival advantage vis-à-vis artificial intelligence?

JOSEF PARVIZI: Oh, hell yes! Without consciousness we would have no feelings of hunger, thirst, sexual desire, fatigue and fear. You tell me if these feelings have no survival advantage? Be my guest, and say no. The whole world will bet against you.

What is consciousness, anyhow?

BILL NEWSOME: It's hard to say what consciousness is when it's so unapproachable from neuroscientists' standard reductionist approach. As philosopher Thomas Nagel has argued: Our contemporary science is inherently third-person, but consciousness is inherently first-person. Anything that takes us away from the first person takes us further away from (not closer to) the primary phenomenon we're trying to understand. So I (along with others) think there is a serious question whether our third-person science can fully understand consciousness.

Can AI become conscious? What would that take?

GIOCOMO: It's possible that our metrics for identifying consciousness may be, at some point, inadequate to truly differentiate between the human brain and AI.

NEWSOME: Whether machines will become conscious is an interesting question, but given how little we know about the physical basis of human consciousness, how can we say anything at all worthwhile about machines? Our digital computers are based on architectures and modes of computing that are totally different from those of biological brains.

PARVIZI: For AI to be potentially conscious as we know it, in other words, for AI to be potentially conscious like a human being, the AI algorithm has to be organized like the human brain, it has to be born and raised in a social environment featuring continuous interaction with others, and it has to reside in a biological organism identical to a human's frail body, with its bladder getting full and urging it to run to the bathroom, or its flesh hurting because it just touched a super-hot surface.

If you re-create a humanlike organism and put a human-brain-like organ in it, there you go: You have a conscious AI. But it is no longer called AI; it is called a cloned human!

Does consciousness require connections to a body?

NEWSOME: I believe consciousness is a biological phenomenon. One thing I feel pretty sure of is that no AI will become conscious in the human or even nonhuman animal sense until it is loaded into mobile bodies and has to work out the really big problems of survival (energy extraction and conservation, competition); reproduction; and the risks, benefits and norms that come along with advanced social cooperation. I think reward, aversion, motivation, etc., are required for the kind of human and to some extent animal consciousnesses that we are familiar with.

Appetites and desires serve fundamental organismal goals and purposes, and I have no idea how a machine can be endowed with these kinds of subjective feelings. Appetite and desire can be simulated, I suppose, as cost functions (or something like that), but does the simulation approximate the real thing, or are the two totally different? Can a machine have goals and purposes other than what is programmed into them by a human?

I don't see how this can happen with the current style of AI, but if AI is indeed loaded onto mobile bodies and faced with solving the problems of survival, perhaps the AI could, over time, modify its human-given motivations and goals and develop its own. Who knows?

Despite my skepticism, I do share the uneasy sense that with AI we are at the threshold of a new, very uncertain age in human and scientific history. While AI may not acquire a human form of consciousness, it may become conscious in ways that are incomprehensible to us. There may be worlds of consciousnesses out there just waiting to be realized in complex, learning beings. Again, who knows?

'THERE MAY BE WORLDS OF CONSCIOUSNESSES OUT THERE JUST WAITING TO BE REALIZED IN COMPLEX, LEARNING BEINGS. WHO KNOWS?'

Can we coexist? Can we all just get along?

SOLTESZ: My only hope for what I admit is my otherwise dark vision of our future is the fact that we as humans have had nukes for decades by now but so far have resisted blowing up the world, so maybe humans will agree to limit AI applications in a practical and effective way.

One other crazy thought: It appears curious to me that we arrive at the dawn of the AI age around the same time as we enter the age of deep-space human travel (plus or minus a few years or maybe a decade). So perhaps humans and AI systems can "agree" right from the start to divide the galaxy into distinct domains of neighborly coexistence. **SM** — *Contact Bruce Goldman at goldmanb@stanford.edu*

MELISSA PICKETT

WHAT'S DRIVING THIS TUMOR

Aiming to stymie breast cancer through gene testing and AI

BY KRISTA CONGER
PHOTOGRAPHS BY LESLIE WILLIAMSON

MELISSA PICKETT didn't expect her breast cancer diagnosis. As a new mom, the 33-year-old college professor had a lot on her mind, including getting to her six-month postpartum checkup in November 2022.

When her doctor found lumps in her breast, Pickett dismissed them as a side effect of breastfeeding. She wasn't worried. She had no family history of breast cancer, and she was used to false alarms. "I have a history of benign breast lumps going back to when I was 16," Pickett said. "So I've had a lot of scans and imaging over the years. It was always nothing."

But this time was different.

When her oncologist at Stanford Medicine, associate professor of medicine Melinda Telli, MD, told Pickett about a newly launched clinical trial for people recently diagnosed with her type of breast cancer, Pickett, who has a PhD in toxicology and is an associate professor of genetics at San Jose State University, didn't hesitate to enroll.

"I've chosen to dedicate my life to the study of basic biology," Pickett said. "Usually I'm the researcher, rather than the subject. But if my cells can be of use, if my case can increase our understanding of the biological processes underlying this disease, I really wanted to participate."

Pickett is no stranger to Stanford; prior to her appointment as a faculty member at San Jose State she was a postdoctoral scholar in Stanford University's biology department studying how epithelial cells (cells that line the inside of organs or glands and that make up the outer surface of the body) orient themselves in three dimensions — designating a region of a cell the top and another as the bottom, each with different functions. Ironically, loss of this orientation, or polarity, is a key hallmark of cancer cells. A few rogue epithelial cells in Pickett's breast had turned against her.

The trial, which is called Terpsichore after the Greek goddess of lyric poetry and dancing, is a novel effort to gather genetic data from a patient's tumor immediately after diagnosis and before treatment has started — and to use that data to attack the biological signals predicted to drive the tumor's growth in its infancy. It wouldn't have been possible without artificial intelligence, which not only enabled the researchers to sort breast tumors into subgroups based on a dizzying array of genetic information but also to home in on shared attributes likely to be sensitive to treatment with existing drugs.

The researchers in the Terpsichore trial will also use AI to track the outcomes of the trial's participants, noting whether and which tumors respond to a two-week window of experimental treatments, as well as helping them follow the patients in the years after their initial diagnosis.

"Most cancer clinical trials focus on metastatic disease, after the standard of care has already failed the patient," Christina Curtis, PhD, Stanford Medicine's director of artificial intelligence and cancer genomics, said. "But by that time, the cancer has had time to mutate and accumulate more genetic changes that drive its growth and make it resistant to treatment."

Focusing on newly diagnosed, untreated tumors will reveal their original "starter pack" of mutations when the cancer is most vulnerable, Curtis and her colleagues reason. They hope this approach will improve breast cancer care for all patients. They also hope it will help answer one of the most pressing questions in breast oncology today: Why do about one-quarter of people with Pickett's type of cancer, categorized as hormone receptor-positive, HER2-negative breast cancer, face a significant risk of recurrence decades after their diagnosis?

That broad window of risk far exceeds the five-year period cited by many oncologists as a reassuring milestone after which a patient can be considered cured. And the likelihood of recurrence — around 50% — for this subset of patients surpasses even that of triple-negative breast cancer, which has fewer treatment options and higher overall mortality than other breast cancer types.

Until recently, doctors had no way of knowing which of their breast cancer patients were at higher risk. But a body of work by Curtis and her colleagues over the past decade has pinpointed genetic changes that can be used to categorize breast cancer types into 11 clinically important groups and identify which of them are at heightened risk of recurrence. They did so by training computers to exhaustively analyze tumors' genomes — the complete set of genetic blueprints encoded by their DNA — and transcriptomes — the genetic messages, or RNA, that hint at the genes and proteins the cancer is using to survive.

THE IDEA IS NOT UNIQUE: Cancer researchers are increasingly turning to AI to parse the exponentially growing amount and types of data now gathered from patients. "This technology, and the power of the computational methods used to analyze this kind of information, will completely revolutionize how we think about the disease process," said Jennifer Caswell-Jin, MD, assistant professor of medicine and principal investigator of Terpsichore. "In the past, we studied sections or samples of tumors; now we can analyze individual cells to identify new drug targets."

They're ready to apply their findings to newly diagnosed people like Pickett, with the aim of targeting the Achilles' heels of tumors earlier in treatment than ever before.

"We believe the changes that increase the risk of recurrence years later are already there in the very earliest cancer cells," Caswell-Jin said.

The researchers hope that the Terpsichore trial will help them reduce the likelihood of recurrence for high-risk patients and that it will reveal the biological underpinnings of how some tumors are able to cool their heels for years, remaining undetected before roaring back to cause a second, devastating round of disease.

BREAST CANCER LINGO CAN BE CONFUSING. Not only are tumors categorized based on where in the breast they occur and which breast tissues are involved, but they are also grouped based on the proteins produced by the cancer cells. The presence or absence of certain proteins, such as human epidermal growth factor or the receptors for the hormones estrogen and progesterone, give clinicians clues about which biological signals are telling the cells to grow.

Blocking those signals with drugs or antibodies can slow or stop a tumor's growth, and clinicians use this kind of molecular profiling to determine whether a breast cancer patient needs chemotherapy, radiation or hormone therapy like estrogen blockers — and for how long. (Most patients will also have some type of surgery to remove the cancerous tissue, either before or after other treatments.)

Broadly speaking, hormone receptor-positive breast cancer responds well to treatment, as do cancers in which the human epidermal growth factor receptor, HER2, is expressed at high levels. Cancers that don't have elevated levels of estrogen receptor, progesterone receptor or HER2 — known as triple-negative breast cancers — are more difficult to treat and more deadly.

The concept of separating breast tumors into categories to guide treatment decisions and prognoses isn't all that new. But at every step, it's been limited by the technology available at the time. The four subgroups described above (hormone receptor-positive, HER2-negative; hormone receptor-positive, HER2-positive; hormone receptor-negative, HER2-positive; and triple negative) are often determined by a cell-staining technique called immunohistochemistry that has been around since the early 1940s. Another test, OncotypeDX, was developed in 2004 and is based on the expression levels of just 21 genes. It is used to determine a patient's five-year risk of recurrence and whether they should receive chemotherapy.

In 2012, Curtis, then at the University of Southern California, led a study that took a more complex approach. The researchers overlaid information about a patient's genome — the DNA inherited from their parents — with that of the DNA sequences and RNA levels found in their cancer cells. RNA mes-

'WHEN THE DATA FROM THOUSANDS OF PATIENTS WAS COMPILED, IT SUDDENLY BECAME VERY CLEAR THAT THERE IS A SUBSET OF PEOPLE AT SIGNIFICANT RISK.'

sages, which are selectively copied from DNA in the cell's nucleus before traveling to its protein-making machinery, provide a snapshot of a cell's operating instructions: divide now, make more of this gene, fire off a chemical signal to a nearby cell, etc.

The researchers took this approach because cancer cells are a messy bunch. The very act of running off the rails — casting aside any semblance of orderly growth or concern about cellular rule breaking — virtually ensures that they bobble the delicate series of events needed to correctly copy and divide their DNA before each cell division. Like a wobbly top, every generation tilts a bit more out of control — accumulating an increasing number of mutations and even adding or losing copies of whole genes willy-nilly. As a result, cancer cells often have variable numbers of copies of important genes in their DNA, a genetic outcome known as copy number variation.

Often, this slow-motion molecular car crash results in the cell's death. But sometimes changes occur that increase the cell's fitness and allow it to climb to the top of the evolutionary dog pile. Cancer biologists call these changes drivers. In theory, blocking them will deal a significant, perhaps fatal, blow to the growing tumor.

The type of multifaceted analysis Curtis and her colleagues were attempting — comparing DNA sequences from healthy cells with DNA sequences and RNA levels from tumors, and doing so for multiple patients — is complicated. Too complicated, in fact, for any one person or laboratory team to tackle. Instead, the researchers fed the information into a computer algorithm in an approach called unsupervised machine learning — allowing the computer to sift through millions of comparisons and derive its own conclusions based on the data available. It's a hands-off approach that avoids bias.

"We wanted to see what kind of groupings would form in the data with minimal supervision," said Curtis, who is now the RZ Cao Professor and a professor of medicine, of genetics and of biomedical data science at the Stanford School of Medicine. "This allowed us to see these cancers through a whole new lens and identify novel subgroups of disease."

When Curtis published her 2012 study identifying the subgroups, doctors didn't know whether or how to use the information to guide treatment. But in 2017, a different group published an eye-opening analysis of 75,000 people diagnosed with hormone receptor-positive, HER2-negative breast cancer that, for the first time, showed that about one-quarter of patients had a 50% chance of their tumors recurring even decades after their initial diagnosis. Unnervingly, even some patients whose cancers had not spread to their lymph nodes (a measure of metastasis) at diagnosis experienced recurrences at much higher rates than had been previously grasped.

CHRISTINA CURTIS,
Stanford Medicine's director of artificial intelligence and cancer genomics

"Doctors had seen unusual, late recurrences before in individual patients, but the patterns had not been systematically analyzed," Curtis said. "When the data from thousands of patients was compiled, it suddenly became very clear that there is a subset of people at significant risk."

In 2019, Curtis, Caswell-Jin and other researchers at Stanford Medicine and the University of Cambridge published a paper in *Nature* describing how it's possible to combine information from immunohistochemistry and their new subgroups — termed integrative clusters — to not only predict which people were at increased risk of late recurrence but also to identify a subset of people with triple-negative tumors who were unlikely to see their cancers return after five years.

The researchers found that four of the 11 integrative subgroups were significantly more likely to return even 10 to 20 years after diagnosis. An analysis of each of their DNA and RNA profiles hints at possible reasons, but the driving factors aren't the same for each group. For example, although each of the four subgroups has tumors with increases in the numbers of copies of several cancer-associated genes called oncogenes, they differ in the number of copies of other genes involved in cancer cell survival and cell proliferation, including a notorious cancer driver called Myc. Many of the genes they appear to rely on for growth are known, and there are already approved drugs that block their actions.

"This molecular profiling gave us the different categories of tumors and helped us understand how these groups fare over time," Caswell-Jin said. "Now we have the information to begin to figure out the right drugs and treatments to interrupt that path to poor outcome."

TERPSICHORE IS AN APT NAME FOR an effort that requires an intricate dance to delicately balance patient care and a multi-armed experiment with many moving parts, each of which needs to mesh seamlessly over a period of about three weeks.

"We need to do this fast enough that we don't delay the standard of care for these patients," Curtis said. Currently, an in-depth genetic analysis of breast cancers is usually conducted only as a last-ditch effort to fight advanced metastatic disease, and the time pressure is less because the patients are already undergoing treatment. For Terpsichore, the researchers have set a goal of nine days in which to gather and analyze the genetic information needed to categorize each patient's tumor into high or typical risks of recurrence using new and improved approaches optimized for clinical samples. Because only one-quarter of people with hormone receptor-positive breast cancers will fall into the high-risk categories, Curtis and Caswell-Jin expect they'll need to screen hundreds of people to find the 150 they'd like to include in the trial. Of those enrolled, about one-third will be in categories predicted to have a typical — that is, low — risk of recurrence after five years, and the other two-thirds will fall into groups with a higher risk.

Once patients enroll, they will be designated at random to receive either standard treatment for their cancers — the control arm of the study — or a 14-day treatment with drugs predicted to block the biological pathways that drive the growth of their tumors — the experimental arm. After two weeks, the researchers will assess the effect, if any, of the treatment on the growth of the tumor. All patients will then undergo a conventional course of treatment, including surgery and a yearslong course of hormone therapy. As the years tick by, Caswell-Jin and Curtis will monitor the participants' health and disease status.

SINCE PICKETT'S DIAGNOSIS, she's had a lumpectomy to remove the cancerous tissue, followed by radiation and hormone therapy to stop the growth of her estrogen receptor-positive cancer cells. She is matter-of-fact about having been randomized to receive standard care, rather than the experimental intervention. "I knew my participation in the trial would not be likely to have any direct benefit to my health," she said. "But without these types of studies, researchers have no way of knowing if they are targeting the right pathways. My prognosis is good, but I know my cancer might return in 20 or even 30 years."

Pickett, who would like to have another child, is particularly concerned about the reproductive effect of cancer drugs. "Cancer drugs have improved so much during the past 50 years," she said. "They used to be incredibly toxic, but now we can target some specific pathways in some cells. Maybe one day we won't have to shut down estrogen production entirely for a person with estrogen receptor-positive cancer. We started with a wrecking ball approach; now we're down to a hammer approach; maybe one day, with studies like these, we can get down to a needle."

Curtis and Caswell-Jin envision a future where even more layers of information can be integrated into ever more sophisticated models of breast cancer biology to help realize Pickett's vision.

"We could integrate the genome and transcriptome data from a tumor with information from pathology, radiology and even spatial data showing where proteins are in a cancer cell or identifying neighborhoods of cell types," Curtis said. "The more data we have, the more powerful these AI approaches can be. We should be leveraging it all to make advances much faster. We have the tools; it's on us. Let's move the dial so patients can benefit as soon as possible." **SM** — *Contact Krista Conger at kristac@stanford.edu*

AI-BOOSTED BIOPSY SCRUTINY

With all patients' tissue samples digitized, AI-assisted analysis begins

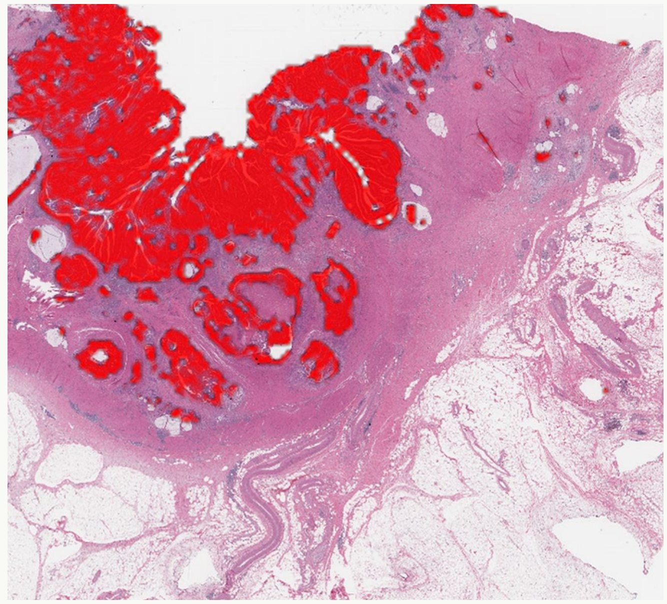

ANALYZING VIA AI

When images of tissue samples are digitized, they can be analyzed via artificial intelligence to identify cancerous areas. Here a pathology image of colon tissue is shown being analyzed by an AI-driven diagnostic program.

ZHI HUANG

FOR OVER 100 YEARS, pathologists have peered through microscopes and called on their medical training to decipher the meaning of a displaced cell or a malformation of tissue structure. Now, artificial intelligence is revolutionizing that practice. Pathologists at Stanford Medicine recently began collecting digital images of tissue samples, and they've designed an AI tool to sift through the mounds of information contained in each pixelated cell and determine its significance.

Stanford Medicine researchers expect the technology to quickly transform patient care.

"Computers can match patterns across millions of images in ways that would be difficult for humans to comprehend," said Thomas Montine, MD, PhD, professor and chair of pathology. "Now, we can begin to unleash the power of artificial intelligence and machine learning techniques onto these data in ways we've only imagined."

Tissue samples of tumors biopsied from patients can tell pathologists what type of cancer they're dealing with, how aggressive it is and whether a treatment appears to be working. Many other diseases and disorders also reveal themselves under the pathologist's trained eye, including infections, genetic disorders and brain diseases.

Stanford Medicine is one of only two academic medical centers in the country with a fully digitized pathology department. The scarcity of such departments reflects the steep price and broad culture change necessary to make this leap. Amassing enough data to train AI to recognize and categorize specific images is also difficult, particularly because many digital images are owned by private health care systems with little incentive to share.

"Pathology around the country still remains mostly manual," Montine said. "A piece of tissue from a medically concerning mass or lesion comes into the pathology lab, where it is cut into thin sections and placed on glass slides and stained with chemical or biological material that will reveal structures and cells. A pathologist looks through a microscope and captures that information in their mind and uses their medical training to interpret a biological meaning. Digitizing this process means that this image is instead captured by a computer and can be easily shared for discussion, teaching and learning."

Although the images, which are stripped of any information that could be used to identify an individual patient, are captured digitally and categorized by the AI software, the pathologist remains firmly embedded in the process.

"We wanted to build an artificial intelligence tool with a human in the loop," Montine said. "So the computer can screen the images, but the final judgment and diagnosis is left to the pathologist."

Training AI to recognize patterns and images is a bit like teaching toddlers their colors by scattering a pile of colored Lego on a kitchen table and pointing out the yellow bricks in various shapes and sizes. After enough repetition — this is a yellow brick, this is a yellow brick, this is a yellow brick — children can pluck out the color themselves. Another round of teaching can lead them to sort large yellow bricks from small yellow bricks, and so on. With pathology images, the process is similar, but far faster and more iterative — homing in not on large yellow Lego bricks, but on breast tissue with hints of ductal carcinoma in situ, for example, or lymphoma cells sneakily hobnobbing with healthy white blood cells.

"This tool has the exciting capacity to extend human capability," said James Zou, PhD, assistant professor of biomedical data science, who partnered with Montine to create the software. "We envision having AI and humans work together as a team to save time and improve patient outcomes. Currently, we rely on clinicians evaluating images by themselves. But there can be tens of thousands of individual cells in one image, and some diseases cause only very subtle changes."

Alone, images are just images. For a computer to learn from them it's necessary to have expert text that describes the image ("yellow brick," "yellow plate," "yellow hinge," for example, to continue with the Lego example). Montine and Zou have launched an effort called Open Pathology to scour the internet for expertly annotated pathology images they can use to further train their tool. They've gathered several million so far. They envision a day when the data they've collected, which include both image and text, are linked to patient information and outcome from electronic medical records, or overlayed with genetic, molecular and spatial data in ways that allow AI to discern ever more complex patterns and associations beyond what even the most dedicated pathologist could perceive.

"This tool will help us improve care for our patients while doing our jobs faster, cheaper and more efficiently," Montine said. "The sky is the limit." **SM** — BY KRISTA CONGER

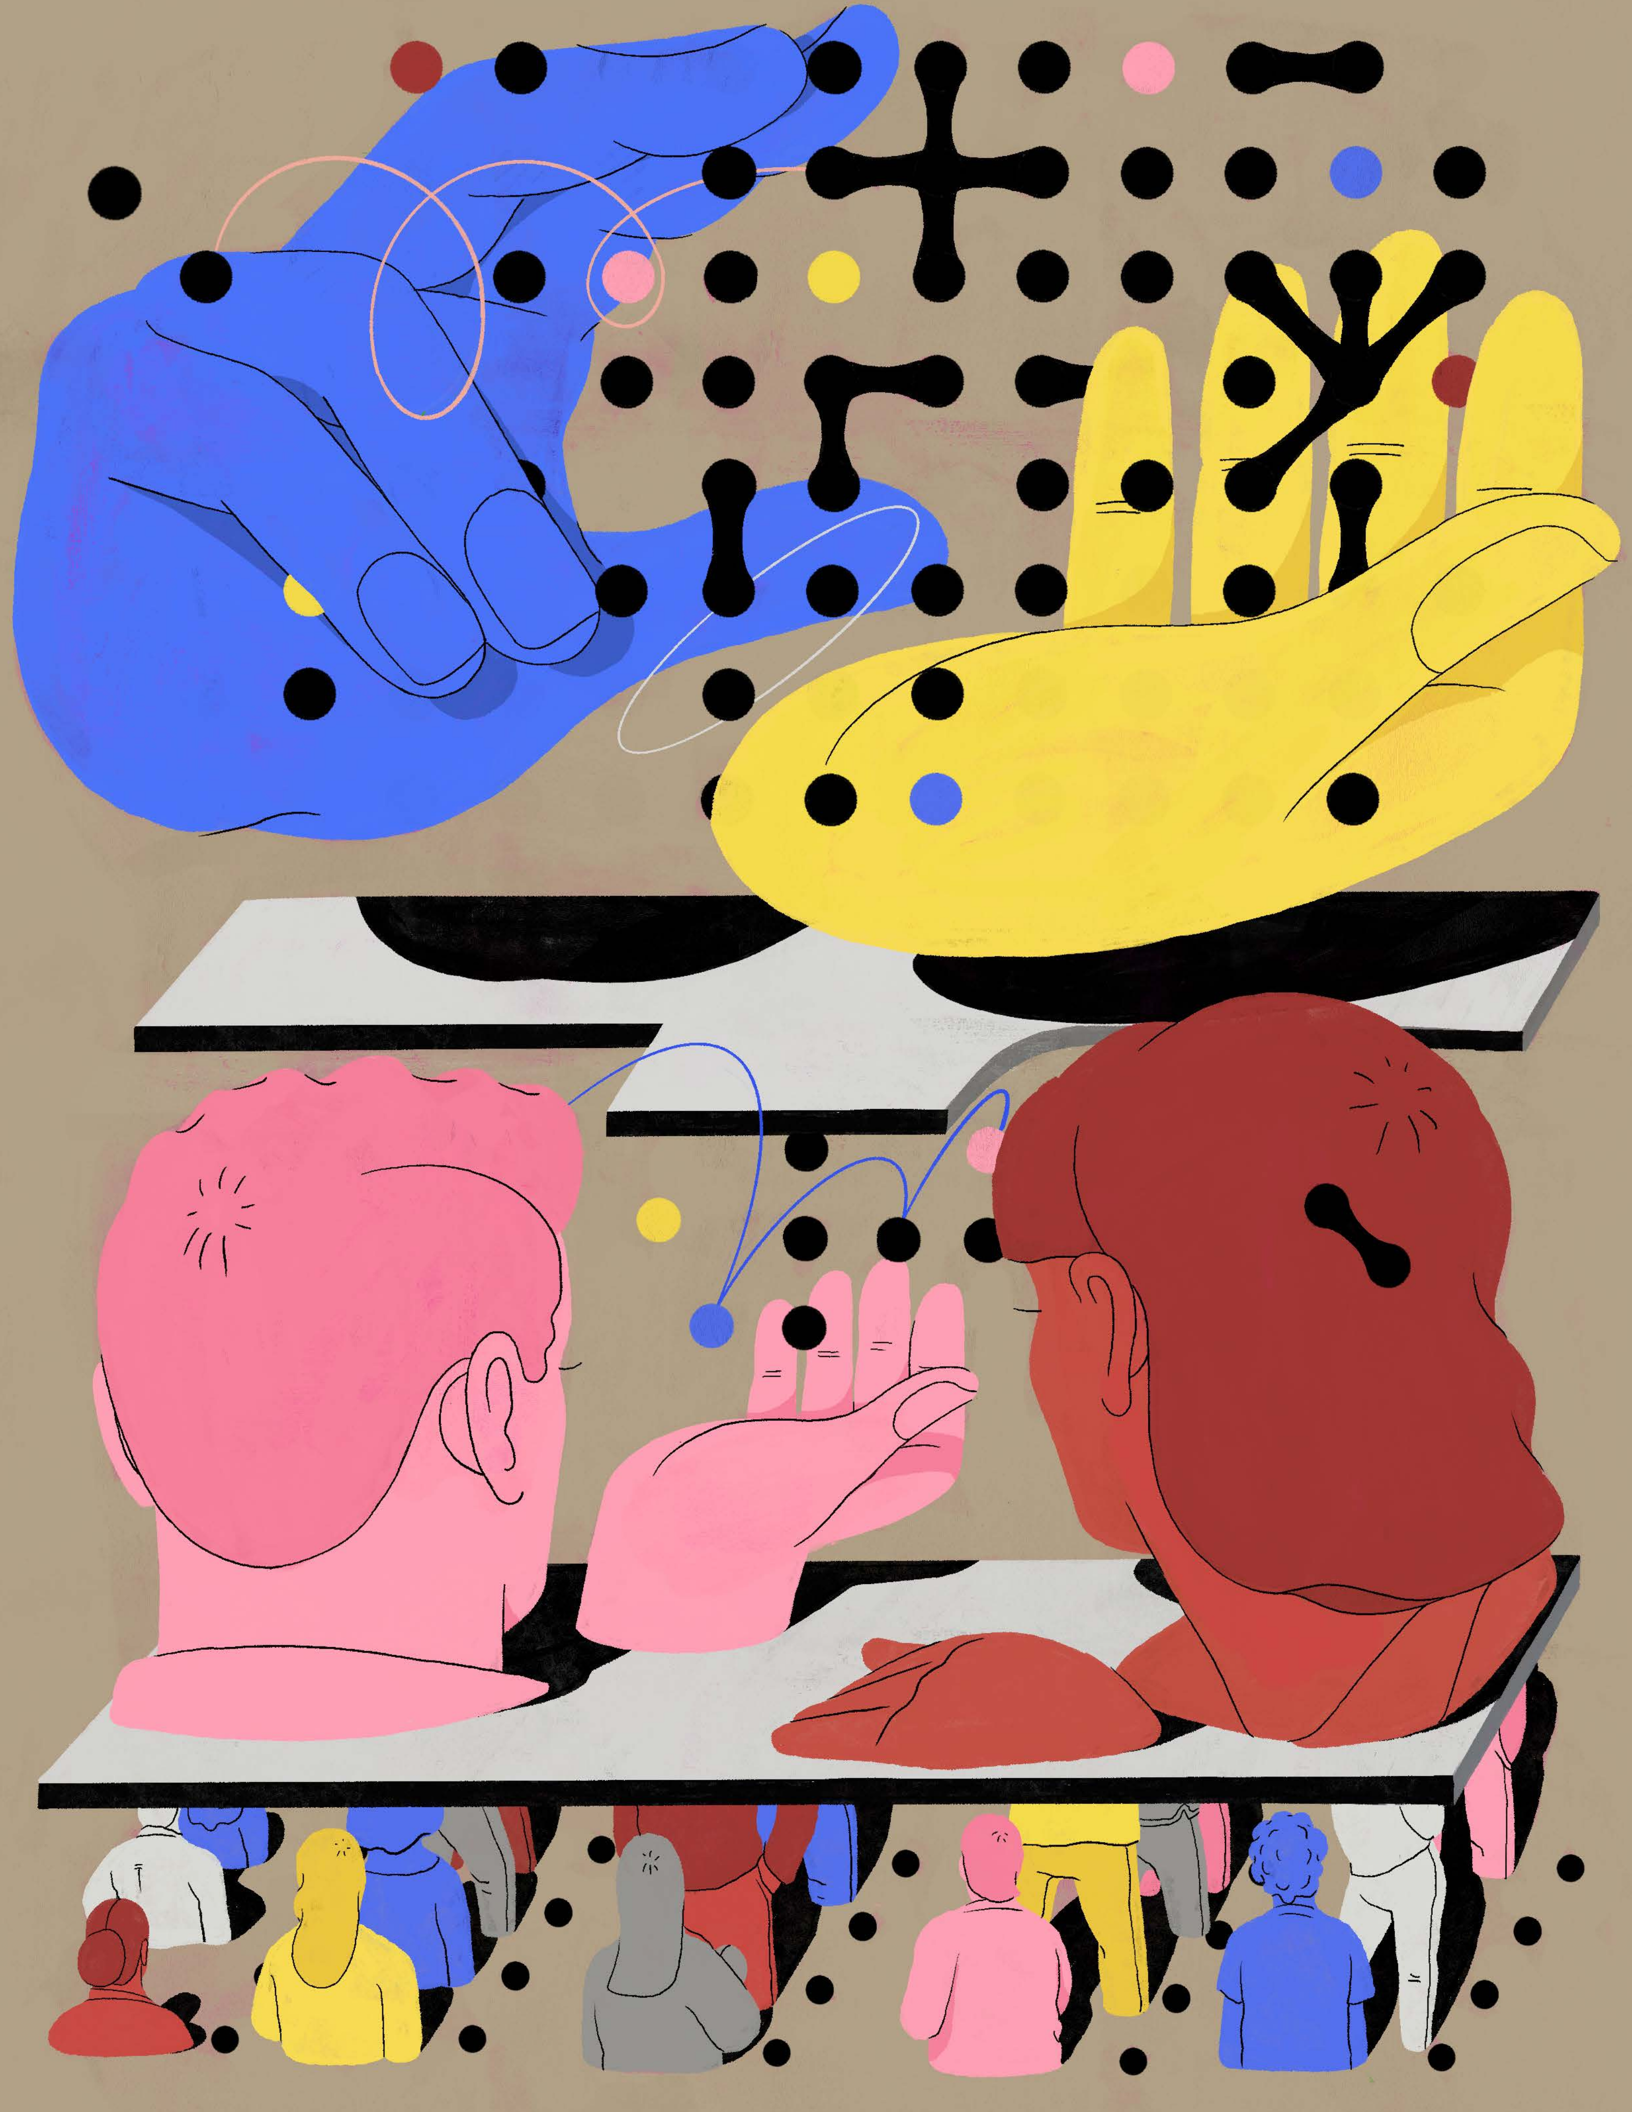

ADDING ETHICS TO THE MIX WHEN DEVELOPING HEALTH CARE AI

A bioethics expert is designing exercises to encourage ethical responsibility in programming

BY SARAH C.P. WILLIAMS

THESE DAYS, IT ISN'T JUST DOCTORS AND NURSES WHO KEEP A HOSPITAL RUNNING SMOOTHLY: It is also computer programs. Data scientists have developed scores of brainy algorithms to pinpoint patients at risk of complications, catch errors in medical records or prescriptions, fast-track paperwork and billing, and even diagnose patients.

While new doctors usually recite some version of the Hippocratic Oath at their medical school graduation — swearing to uphold ethical standards in treating their patients — programmers who develop AI for the health care industry are rarely given formal ethics training. Mildred Cho, PhD, a professor of pediatrics and associate director of the Stanford Center for Biomedical Ethics, is trying to change that.

ILLUSTRATION BY BRYCE WYMER

"Developers are often not from a medical background and haven't spent years thinking about this moral framework — how things like respect and justice and personal principles spill over into medicine," Cho said. "But as we start seeing artificial intelligence programs being used more widely in medicine, it's important that developers think about the real-world ethical implications of their work."

Over the past five years, Cho has interviewed dozens of programmers who work in a variety of settings, all creating health care-related machine learning programs. With machine learning, developers input existing data on patients into a computer, which pinpoints patterns that might not be obvious to a person. Using these patterns, the machine learning algorithm can then analyze new data from outside the original set. Machine learning algorithms can be used to identify patients who are at risk of dangerous complications like malnutrition, falling or infections, for instance, and flag them for additional attention or treatments.

"Almost anything you can think of in medicine is being tackled with machine learning right now, because so much of medicine is about pattern recognition," Cho said.

But artificial intelligence algorithms also pose hazards: They can misdiagnose patients, fail to identify people at risk of complications, or reveal pieces of private information.

AI can also exacerbate existing biases in the health care system: If doctors are less likely to diagnose women or minorities with a particular condition, machine learning platforms will assume that people in those groups develop the condition less often, perpetuating the bias. Scientists at Duke University Hospital, for instance, designed an AI program to identify children at risk of sepsis, a dangerous response to an infection. But the program took longer to flag Latino kids than white kids, possibly delaying the identification and treatment of Latino children with sepsis. The bias, it turned out, existed because doctors themselves took longer to diagnose sepsis in Latino kids. This taught the AI program that these children might develop sepsis more slowly or less often than white children.

"What's really lacking in AI right now are standards for evaluating data quality," Cho said. "What does it mean to have a safe and effective AI tool in a health care setting?"

'ALMOST ANYTHING YOU CAN THINK OF IN MEDICINE IS BEING TACKLED WITH MACHINE LEARNING RIGHT NOW, BECAUSE SO MUCH OF MEDICINE IS ABOUT PATTERN RECOGNITION.'

WHEN CHO INTERVIEWED DEVELOPERS, she was surprised by how many admitted the potential pitfalls of their products; she had suspected they might not be aware of all the risks and biases.

"Despite not having training in medical research, most developers were actually able to identify quite a few potential harms that might come about as a result of their work," she said. "They were thinking on a much bigger level than I thought they might be."

But when she asked them what to do about these potential downsides, the developers tended to pass the buck. They said that someone else — their bosses, their companies, the health care systems using their products, or physicians themselves — should be making sure the AI programs were used ethically and responsibly.

"The phrase I heard the most often was 'at the end of the day,'" recalled Cho. "They would shrug and say things like, 'At the end of the day, this is a business' or, 'At the end of the day, I'm just a low-level data scientist and it isn't my problem.'"

Cho doesn't agree. She wants to teach AI developers that small decisions they make while coding can have massive implications for patient care. So in 2022, she began a pilot program offering two-hour group training sessions in ethics for AI programmers. In each session, she asked the developers to begin brainstorming what they would need to make a machine learning algorithm that predicted diabetes risk.

At first, she told them they were making a research tool and asked what they'd need to consider in creating it. Their list, she said, was mostly technical: They needed high-quality patient data and good existing models of what health and demographic factors influence diabetes. Next, Cho asked them to repeat the exercise but to assume their program would be used in a large health care system rather than only for research. At this point, she said, the developers started talking about clinicians for the first time, imagining how doctors and nurses might implement the AI into their practice. These kinds of considerations, Cho explained, can ultimately change how the AI is designed in the first place.

Finally, Cho and her colleagues asked the developers to imagine that they were creating the AI tool for diabetes screening not just for any health care system, but for their own health care system. Suddenly, the developers began talking about the patient perspective of the algorithm, discussing topics like how to ensure that patient privacy is maintained and that health care remains high-quality.

"They actually switched their entire perspective and considered completely new aspects of the project," Cho said.

Her hope is that developers who go through this exercise can apply the lessons to their work, putting themselves in clinicians' and patients' shoes while creating health care-related AI programs.

"What I want is for developers to move toward thinking about their own ethical responsibilities, anticipating what harms their programs could have, and pulling those ideas into the design phases of their work," she said.

So far, Cho has tested the training with five groups of four developers. Eventually, she'd like to try it in a workplace environment, with developers carrying out the exercise with real AI software that they're in the process of coding, rather than a hypothetical diabetes-prediction tool.

Peter Washington, PhD, an assistant professor of information and computer sciences at the University of Hawaii, participated in Cho's pilot program when he was a graduate student at Stanford University. Washington has built machine learning programs to detect autism and has worked as an intern at Google, Amazon and Microsoft Research. He now leads a digital health research lab in Hawaii building a variety of machine learning models for diagnosis and disease tracking. He said that programs like Cho's, which encourage developers to think about the applications of their work, can help improve the privacy and fairness of AI. Now, he integrates ethics lessons that he learned while interacting with Cho and other members of the Stanford Center for Biomedical Ethics into the computer science classes he teaches.

"Ethics is not usually taught in computer science programs, and if it is, it's an elective rather than a required course," Washington said. "But I think it's really important for developers to understand core sociotechnical issues like the inherent tradeoffs that exist between things such as privacy and accuracy."

In other words, the more patient details an AI program has access to, the more accurate it may be but the more likely it is to invade privacy. He said it's especially important for developers to think deeply about what data they are using to train their machine learning programs. A program that has learned patterns from patients who are mostly white, for instance, may not work as well for Black or Latino patients; a program trained on data from a small, rural hospital may not draw accurate conclusions if used in a large, urban hospital.

'ETHICS IS NOT USUALLY TAUGHT IN COMPUTER SCIENCE PROGRAMS, AND IF IT IS, IT'S AN ELECTIVE RATHER THAN A REQUIRED COURSE.'

Even when they don't have ultimate control over when and how their products are used, AI developers can make changes that go a long way towards solving these challenges, Washington said. For instance, fairness metrics — numbers showing the potential biases within an AI model — can add transparency to an AI program and help users understand how the model might perform differently in different settings and for different populations of patients.

"You can write a few quick lines of code that calculate fairness metrics that will uncover potential biases in the model," he explained. "This is incredibly easy to implement if you have the demographic data available, but very few groups are incorporating this as standard practice."

Ultimately, the companies developing AI tools and the health care systems deploying them in hospitals and clinics do have to take responsibility for the ethical use of AI in medicine, Cho said.

"It is hard to completely pin responsibility on developers for how things are used after they're released," she admitted.

But the more developers think about how to minimize bias in their tools and be transparent in the strengths, weaknesses and best applications of their AI tools, the easier it will be to use their products in appropriate — and ethical — ways, she said. **SM**

— Contact Sarah C.P. Williams at medmag@stanford.edu

and yet, there's hope

A family devastated by cancer plants seeds to vanquish the disease

BY HANAE ARMITAGE

Anyone who knows Aruna Gambhir would never call her fragile. In her lifetime she's earned three degrees, including an MBA; nursed family members through multiple cancer diagnoses; led a biotechnology company for more than a decade; and survived several health threats of her own. But on a recent quiet, late summer afternoon, as she leans gently against the back of a plush armchair, her petite frame dwarfed by the oversized cushions, her physical state is unmistakably delicate.

She has cancer. At 61, she's fighting it for the third time.

Since November 2021, when doctors detected Gambhir's latest bout of the disease, she has waged a biological battle against the condition, myelodysplastic syndrome, in which early forms of blood cells fail to mature into healthy cells: A bone marrow transplant wiped out and replaced faulty blood cells, and chemotherapy killed off any remaining cancerous cells. Despite all the treatments, her illness persisted, leaving her health in a tenuous balance.

While sitting in the light-flooded living room of her Portola Valley home, Gambhir tells her story. Cancer has evoked her family's darkest days. She has lost her father, her brother, her son and her husband to the disease. She has assumed the role of a caregiver, even when it seemed daunting or unfair, even as she struggled for her own life.

Still, she holds onto hope — and it's born from cells of the very tumor that took her son's life, which are now being studied in a lab at Stanford Medicine. Gambhir hopes her family's story will expedite progress in cancer research by bringing together government, biotechnology companies, and public and private research funders.

Her ultimate hope is that one day the biological perils that befell her family will be tamed and cancer will be transformed from a fatal illness to a disease that can be bested.

A biological secret

IN 1982, when Gambhir was 22, her relationship with her father, Satya Bodapati, took an unexpected turn. He had been diagnosed with esophageal cancer, so, just after she graduated from the University of California, Berkeley, Gambhir starting helping her mother care for him in Cupertino, California. Her father underwent surgery and radiation, and doctors deemed the treatment successful, noting that if he made it to the five-year mark without complications, he would be out of the woods. He did. "But, of course, right after the five-year mark, it came back," Gambhir recalled. Shortly after the recurrence, he died.

TIMOTHY ARCHIBALD

ARUNA AND SAM GAMBHIR AT HOME IN 2016. RESEARCH ON CELLS FROM THEIR SON'S BRAIN TUMOR HAVE LED TO IMPORTANT DISCOVERIES ABOUT CANCER.

She didn't know it at the time, but it was her first exposure to what would become an unwelcome theme in her life: a seemingly endless recurrence of cancer in her family. Gambhir suffered two breast cancer diagnoses, one when she was 37 and another at 47. Then, in 2013, she and her husband, Sam Gambhir, MD, PhD, at the time professor and the chair of the radiology department at Stanford Medicine, were contending with their worst nightmare: Their son, Milan, had been diagnosed with one of the deadliest brain tumors, glioblastoma.

The early and aggressive cancers rose a red flag for their doctors, and the family was tested for and found to have a rare, inherited condition known as Li-Fraumeni syndrome, which essentially turns cancer into a trick candle.

First identified in 1969, Li-Fraumeni syndrome stems from a mutation in the gene p53, known for its role in tumor suppression. A functional p53 gene helps regulate cell growth and prompt repair to damaged DNA, among other functions, which keeps cells and tissue in tip-top shape. But when p53 is hampered, cells with damaged DNA stay damaged, and cell proliferation runs unchecked, drastically raising the risk for cancer. In other words, in people who have Li-Fraumeni, cancer can be extinguished, but its propensity to return is unrelenting.

A patient and a caregiver

GAMBHIR'S FIRST cancer diagnosis came when Milan was a year old. Ten years later, she found out she had breast cancer for a second time. Her doctors caught it early and successfully removed it. Because Li-Fraumeni syndrome put her at serious risk for recurrence — and she had a gravely ill son — she eventually decided to have a double mastectomy and a full hysterectomy, a procedure in which the ovaries and uterus are surgically removed.

"When Milan was in a steady state, I took a few days to go into the hospital and get it taken care of. I didn't want to be sick at the same time as Milan; I couldn't care for him if I got sick again," she said. "It's one thing when a single person in the family is sick; it's another when multiple people are ill. You have to pull together and essentially triage to care for the sickest person."

For nearly two years, Aruna and Sam made Herculean efforts to save their son, bringing together a team of doctors and scientists at the top of their field to find a treatment that could help Milan. "But it wasn't in time," said Aruna Gambhir. Milan died at 16 on May 2, 2015.

Three years later, in 2018, the family was dealt another devastating blow, landing Aruna Gambhir back in the role of caregiver, this time for her husband. Sam Gambhir had been diagnosed with can-

HIS CELL CULTURE BECAME ONE THAT THE MONJE LAB STUDIED OFTEN. 'HIS CELLS WERE A KEY ENABLER OF WHAT I THINK HAVE BEEN OUR MOST IMPORTANT DISCOVERIES OVER THE LAST 10 YEARS.'

cer. "I remember thinking, 'No, how can it be? I'm the one with the mutation, I'm the one who should be next. Why is this happening?"

Sam Gambhir spent his career devising new technologies that could reveal signs of early diseases, particularly cancer. With an approach that balanced innovation and practicality, he had an array of curious and creative technologies brewing in his lab. From a cancer-detecting smart toilet, to an "immunodiagnostics" technique, which repurposes the body's immune cells to report on the presence of disease, his dedication to understanding and flagging signs of tumors earned him a reputation as a pioneer of early cancer detection.

Again, the Gambhirs fought with everything they had, rallying teams of world-renowned experts. He endured different assortments of drug cocktails. "We tried all sorts of innovative things, even a clinical trial of one, but the side effects left him in horrendous pain," Gambhir said.

It was too much, and not enough. On July 18, 2020, seven years to the day of Milan's tumor diagnosis, Sam Gambhir, beloved and venerated by his family and colleagues, died.

"To go through the experience of taking care of and losing your son — that's the first sword through the heart. But to have to do so again, this time for your spouse, who everyone just loves — it was terrible," she said.

Four months after her husband's death, Gambhir learned of her third cancer diagnosis and returned to Stanford Hospital for care. While there, she had a visitor who ignited the first feeling of hope she'd had in a long time. Michelle Monje, MD, PhD, professor of neurology, came to tell her that less than a mile away in her Stanford Medicine lab, a new potential treatment for glioblastoma was growing in a petri dish. And it was thanks to Milan.

A devastating diagnosis

UNTIL HE WAS 14, Milan was the picture of health. Remarkably bright by any measure, with an impressive knack for science and math, he excelled in school, played guitar and spent his high school summers as an intern in Stanford Medicine research labs.

Gambhir remembers the day all of that changed. Milan was leaving the house when the water bottle he was holding slipped from his grip. "I didn't think anything of it at the time, but that might have been the first indication that something wasn't quite right," she said.

That night, something more alarming happened. While eating dinner at the kitchen table, Milan started speaking incoherently. "He was talking gibberish — I thought he was joking around and I told him to stop," Gambhir said. He eventually did, but when he recounted the experience to his mom, he said he couldn't control it. "I was worried it was a stroke," she said.

The Gambhirs took Milan to the emergency department, where doctors conducted a CT scan and an MRI. Sam Gambhir watched the image materialize on the monitor. "He knew exactly what it meant. He knew, in that moment, what was in store for our family — for Milan," she said. Glioblastoma is one of the rarest and most aggressive forms of cancer, and its survival rate is dismally low. Most patients survive only about 14 months.

Milan survived 21 months. During that time, the Gambhir family tried desperately to rid Milan's brain of the mass that threatened his life. Milan had two brain surgeries; underwent chemotherapy and radiation therapy; and received other treatments including withaferin A, a molecule isolated from the plant ashwagandha, an ingredient in ayurvedic medicine.

As a father, doctor, scientist and caregiver, Sam Gambhir was laser focused on finding a solution for Milan. He rallied a network of cancer doctors, some to provide direct care to Milan, and others, including Monje, to conduct research that involved obtaining samples of Milan's brain tumor post-surgery and cultivating its cells in the lab to create glioblastoma tumor models.

Models grown from patients' cells offer the best way to study tumors — they allow scientists to better understand the disease's biological roots and devise targeted treatments. For Milan, Monje led that work, a task she usually delegates to others in her lab. "It was so important to get it right; I did it with my own hands," she said.

From culture to clinical trial

WITHOUT identifying Milan as the donor, Monje shared his cells among researchers seeking clearer understandings of pediatric brain tumors and how to treat them. His cell culture became one that the Monje lab studied often. "His cells were a key enabler of what I think have been our most important discoveries over the last 10 years," Monje said.

In particular, Monje and her lab showed how normal brain function, via neuron signaling, actually stimulates tumor growth.

A protein called neuroligin-3, which, under healthy circumstances, helps regulate neurons' formation of new synapses, was a key culprit in fueling the growth of high-grade gliomas, a class of brain tumors that includes glioblastoma.

Through follow-up studies, Monje's lab showed that, in mouse models of the disease, inhibiting neuroligin-3 halted growth of the tumor for several months. Between 2015 and 2019, Monje and her team continued to elucidate the underpinnings of neuron-tumor signaling and growth. Now, her team is in the early stages of conducting a clinical trial for a drug that could block neuroligin-3 in patients with the same type of tumor Milan had, as well as others like it.

Blocking neuroligin-3 alone won't cure the cancer, Monje said, but it's a big step. The five-year survival rate for glioblastoma is less than 1%, and if growth can be stalled early on, it could open the door for other treatments to attack the mass.

"Michelle came to me and explained the positive things that have come out of Milan's cells," said Gambhir. "It's pretty amazing. She's amazing. If Milan had to die — that something as promising as this could happen feels meaningful."

"All of those studies were based on Milan's cells," Monje said. "My laboratory and Milan's story have been so intertwined. He's really contributed to our work."

And he has inspired others to enter the field.

A powerful inspiration

IN 2017, Michael Quezada, then a Stanford University freshman, had been researching breast cancer in a lab that happened to be one floor below Monje's. He'd been wanting to switch his focus to brain tumors and asked Monje if he could join her team. "I'd known Milan and the Gambhirs since I was in preschool," Quezada said. "He was my childhood best friend." Witnessing Milan's fight with cancer ignited a drive in Quezada to understand — and one day stop — the disease that took his friend's life.

But he didn't share that with Monje when he asked to join her lab. What's more, Quezada had no idea that Monje had personally handled samples of Milan's tumor.

Monje welcomed him into her lab. For three years, Quezada unknowingly worked with Milan's cells, researching glioblastoma and other similar brain cancers. All the while, Monje mentored Milan's best friend — and neither was the wiser.

That is, until, Quezada noticed a detail in a story, "And yet, you try," published in *Stanford Medicine* magazine, which revealed Monje's role in cultivating Milan's cells. "I was in disbelief — it was honestly shocking," he said.

As a high schooler with a terminally ill friend, Quezada remembers feeling helpless. "The only thing I could do was spend time with Milan and be there for him," he said. "I felt like I was at the whim of the tumor. We all kind of were."

After Milan's death, Quezada thought about how he could have helped Milan more or done better in some way. "I constantly reflected on that," he said. "I think when I found out that his cells were actually in the lab that I was working in, and that I had been working with his cells potentially for years, it let me feel like I was doing something — that I did something that was meaningful. It made me feel like I was a little bit less helpless."

But it wasn't until Quezada left to pursue an MD-PhD at Harvard that he finally told Monje of their shared connection. He had written a personal statement for his medical school application, and he shared the letter with her.

Quezada wrote about how watching Milan battle cancer set him on a mission, and it continues to be his driving force today — "to work at the forefront of patient care and the frontier of biomedical research, where we can make advancements so that fewer people find their dreams cut short."

Quezada doesn't often mention Milan's death as his motivation to pursue medicine. It's something he generally keeps private. But he's glad he shared it with Monje. "I hope it meant something to her," he said. "It meant something to me."

A legacy that lives on

MILAN ALWAYS wanted to go to Stanford, Gambhir said. But now Milan's name will be forever linked to Monje's career through an endowed professorship, one of the highest honors given by the university. This fall she was named the Milan Gambhir Professor of Pediatric Neuro-Oncology.

"I'm enormously grateful," Monje said. "It's truly an honor, and there's no name that I would rather carry forward through my work."

"Now his name can live on in a different way," Gambhir said. "It's full circle."

If her family's story signals anything, it's the need for a collective push to move the field forward, Gambhir said. "It will take all constituents. I'm the wife of a big-time researcher who couldn't have done what he did without the private donations and support of his work alongside government grants. They all took a chance on his weird ideas — and those weird ideas became mainstream and changed the field." Sam Gambhir's ideas created a burgeoning field: early cancer detection and diagnostics. Now his wife carries the torch. "To one day be able to detect cancer so early and reliably that it's, at worst, a chronic disease — that will be a big achievement." **SM**

— Contact Hanae Armitage at harmitag@stanford.edu

an act of love

For a family with roots in Asia and a loved one in the hospital, food is more than just food

BY ADELA WU

ILLUSTRATION BY JUAN BERNABEU

My mother had just one special request. When my dad was admitted to the hospital for the final time, she implored his doctors not to remove his nasogastric tube. The tube in question, a thin, yellow, silicon catheter delicately taped to his nose, snaked inconspicuously from his left nostril to drape behind his shoulder. Its purpose was to provide nutrition for him, like other patients who cannot easily swallow solid food.

But I knew, as a neurosurgery resident physician, that the tube's innocuous appearance is deceptive. During its insertion, patients gulp as someone pushes the unnatural-tasting catheter through the nose and nasopharynx and down the throat for over 20 centimeters until the end reaches the stomach. It's not necessarily painful to have a nasogastric tube, but it is an uncomfortable experience.

His oncologist gently asked my mom, "Why do you want Mr. Wu to have this? He doesn't need the tube any longer."

"Please, let him keep it," she pleaded. "If you're telling us that he has no treatments left for his cancer, he needs the tube." In bed before us, my father slept through all these discussions, his skin ashen and taut around his sunken cheeks. My eyes stung with tears. I contemplated how my dying dad's fate, along with my heart, was pulled in different directions while my distressed mom struggled to accept his well-meaning doctors' recommendations.

Then, my mother insisted, "In my belief, he needs food for a full stomach to pass peacefully into the next life." She simply did not want to see my dad suffer in this life ... or the afterlife.

For 17 years Dad lived with a malignant brain tumor, diagnosed when I was in seventh grade. He held the record at his hospital for the most brain surgeries to remove his recurrent meningioma but had always recovered quickly from his procedures, returning to work and his daily walks. Sometimes, when I went home for vacations during college and medical school, I joined his evening strolls around our neighborhood, when the hot and dry Southern California air turned into a pleasant, balmy breeze. He liked to ask me about school and whether I had tried any new recipes. With chagrin, I rambled about a beef (cooked to a chewy char) and broccoli (boiled in a pot I then repurposed as a sauté pan) stir-fry dish I cobbled together with knockoff oyster sauce from Safeway.

While he was a pharmacist by trade, my father's first job after he immigrated to the United States from Taiwan was as

a fry cook, presiding over vats of oil into which he dipped dough sticks to make delicious, golden, fried Chinese crullers. Even after he opened his own pharmacy, leaving his restaurant life behind, Dad always kept several thick cookbooks on his bedside table. They were all written in Chinese, so I couldn't understand the recipes. But every page was splashed with high-definition color photographs of the most delectable foods — red braised pork belly, egg drop soup, crispy roast duck. On Sunday mornings, he whipped up these gourmet dishes and any of my favorite traditional Taiwanese cuisine, made to order.

My dad was not the only chef in our household. Whenever I caught colds, I remember Mom cooking and brewing broths with specialty Chinese goods and medicinal herbs. Now, she does the same for my daughter. Food was her way of showing love to her family. Love was a bowl of crisp and sweet sliced Asian pear with the peel carefully removed. Love, too, was a feast of a dozen home-cooked dishes when I flew home. Love was my mother always asking me, "Have you eaten?" because, in her mind, providing nourishment was the purest form of caring. Not only to my parents but also to others raised in Chinese and many other cultures, food is a gift and an act of love.

My father kept his eyes closed when my mother and I entered the hospital room. When he did open them, his eyelids fluttered lightly and erratically as if he were trying to wake from a bad dream. "Dad?" I whispered.

His doctors painted a grim picture. Because a large tumor was now plugging the normal plumbing of my father's brain, fluid accumulated and caused swelling around his brainstem and throughout his brain. The cancer was growing in a place and at a pace that threatened to compromise his ability to speak, to move his eyes, to swallow and, eventually, even to breathe. The surgeons declined to intervene, saying that any procedure would be too risky. The oncologists had no medications left to offer. The palliative care physicians tried to explain hospice to my mom.

Of course, we couldn't subject my dad, mute and unconscious toward the end of his life, to unnecessary treatments and devices. The doctors pulled his nasogastric tube after several more conversations with my mother. Finally, she had understood Dad was suffering even though he could not say so. In grief, Mom rinsed my dad's mouth every day with foam swabs soaked in fragrant ginseng tea to give him comfort and to restore his "chi" — energy — for three weeks until he died. In her own way, she chose how to provide my dad sustenance, propriety and love in accordance with her own Buddhist and Chinese beliefs.

My mother's fixation on the nasogastric tube and feeding my father was

CONTINUES ON PAGE 57

clinical trials in a dish

Fast-forwarding drug development

BY KRISTA CONGER

In a series of rooms in the heart of the Stanford Cardiovascular Institute, incubators the size and shape of a dorm-room refrigerator hum quietly. Inside each, a surprise: hundreds of people.

Well, hundreds of people's cells, that is.

The incubators house a multitude of palm-sized, clear plastic trays dotted with circular wells, which are filled with liquid ranging from pink to yellowish. Each well contains languidly beating heart cells that genetically match one person who has donated them for research. One person per row; 96 wells per plate; 10 to 20 plates per humid, warm incubator shelf; two to four incubators per room.

For the past decade, researchers in the laboratory of institute director Joseph C. Wu, MD, PhD, have been using these trays of cells to investigate the molecular causes of common heart disorders and test the effects of newly designed drugs on heart and blood vessel cells. They envision a future where the tray-based screening — a technique that Wu, the Simon H. Stertzer, MD, Professor and professor of medicine and of radiology, calls "clinical trials in a dish" — reduces the need for large-scale, expensive and time-consuming experiments on humans and laboratory animals.

Like boxy time machines, the incubators, and their contents, are poised to drastically fast-forward drug development by making it much quicker, less expensive and more precise than current methods.

Clinical trials in a dish may also reduce racial and ethnic disparities that often plague real-world clinical trial enrollment, streamline drug development that's aided by artificial intelligence and even predict which of several possible drug treatments would work best for an individual — the ultimate in precision health.

It's a tall order, but the vision is timely. Late last year, Congress passed the Food and Drug Modernization Act 2.0, authorizing drug developers to seek alternatives — among them, cell-based tests like Wu's — to animal testing, which is now the gold standard when first determining whether potential drugs are effective and safe. The act is a nod to the growing realization that dish-based drug screening of human cells may offer a better way to identify promising new treatments for many illnesses, from cancer to heart disease.

There's certainly room for improvement. It's estimated that the cost of bringing a new drug from the earliest laboratory experiments to a final approval by the Food and Drug Administration hovers between $1 billion and $2 billion. The cost is high, in part because failure is baked into the system; fewer than 10% of drugs that succeed in animal trials go on to be approved for use in humans. And none of them are reliably effective for everyone.

"Right now, much of clinical medicine is based on trial and error," said Wu, who is also the president of the American

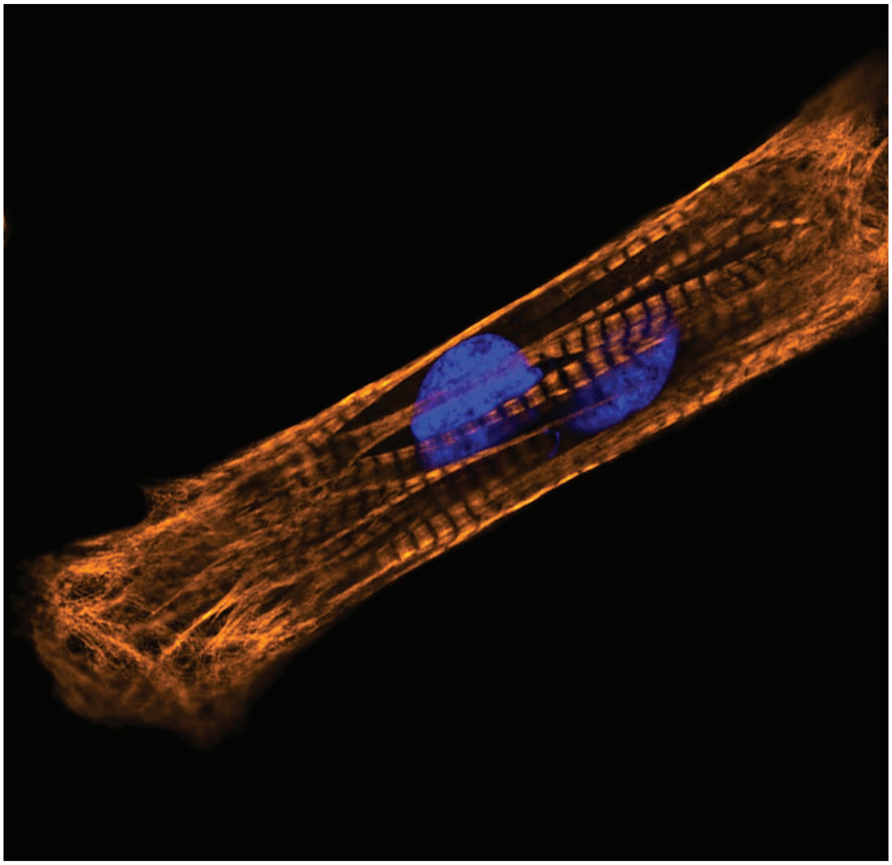

TESTING DRUGS FIRST ON LAB-GROWN HEART CELLS
Researchers have reprogrammed skin cells to create cardiomyocytes, the cells that enable heart muscle to contract, and have begun using them for dish-based drug screening. In this micrograph of a cardiomyocyte, the orange indicates the protein troponin T, while the blue depicts the nucleus.

PHOTO BY JOSEPH WU AND LU REN

Heart Association. "No one drug is effective for everyone. People think because a drug is approved by the Food and Drug Administration it will work for them. But this is not always true."

What exactly is a clinical trial?

CLINICAL TRIALS come in many flavors, including double-blind, randomized and placebo-controlled. There can be multiple groups, or arms, to test various medical interventions, and outcomes may be assessed at multiple time points or triggered by certain clinical readouts. But at their core they are all basically the same: People with the same disease or condition are separated into two or more groups, some of whom will be given an experimental intervention and others who will receive either the current standard of care or no treatment. Doctors then assess which group fares better.

Such commonsense comparisons have a long history, stretching back thousands of years to biblical times. But the father of modern clinical trials is considered by many to be British naval surgeon James Lind, who in 1747 tested six possible daily interventions on 12 sailors with scurvy: seawater; cider; a fiery concoction of horseradish, mustard and garlic (yum!); an elixir of dilute sulfuric acid in alcohol; vinegar; and oranges and a lemon.

After about a week, the two sailors who had received citrus were in markedly better health than their peers. When the British navy began incorporating the routine distribution of citrus fruits and lemon juice to sailors, they virtually eradicated scurvy in their ranks.

In 1938, Congress passed the Food, Drug and Cosmetic Act, which mandated that drugs be proven safe before they can be approved and marketed for routine clinical use. First, laboratory experiments must show the molecule behaves as expected — that it binds to a protein important in disease, for example, or rectifies a disease-associated problem in cellular signaling or function. Often these first rounds of testing incorporate what are known as cell lines, or common, well-characterized lineages of cells grown for years, sometimes decades, in a lab. These predictable, genetically identical cells are the bread and butter of cell-based laboratory studies, but the bread is white and all the slices are exactly the same.

Next, testing in laboratory animals is required to determine that the molecule is safe and, ideally, that it has the predicted physiological effect, although not every human disease can be accurately mimicked in animals. Like the Wonder Bread-like cell lines, however, many laboratory animals are highly inbred and genetically similar. This similarity ensures that the cells, and animals, behave biologically in predictable ways and that experimental results are reproducible over time and among different laboratory groups.

If things go well with the animal testing, a stepwise series of studies, or clinical trials, in humans are then conducted over months or years to further confirm the drug candidate's safety and effectiveness. Molecules that pass these hurdles — about 1 out of every 4 — are tested in a phase 3 clinical trial, which is often conducted at multiple institutions and enrolls hundreds or thousands of genetically diverse participants. These trials are meant to ferret out rare or unusual side effects that wouldn't be obvious in a smaller, more homogeneous group. If the drug remains safe and effective, the manufacturer can apply to the FDA for approval.

Only about 3 in 10 drugs that enter phase 3 trials reach this milestone.

It's a far cry from handing lemons to a scurvy-afflicted sailor, and the threshold for success can vary.

"There's a phrase in the clinical trial world: number needed to treat, or NNT," Wu said. "This refers to the number of people that need to be treated with a particular intervention or treatment to prevent one bad outcome — death, or stroke or whatever is being measured by the trial. People mistakenly think because a drug is approved by the FDA, it will always work for them. But for many drugs, the NNT may be 20 or 30 or even higher. So if the NNT is 30, that means 30 people need to take the drug for one person to benefit from it. And right now we simply don't know which drug will work best for which individual."

Clinical trials in a dish stand to upend this system. The first rounds of testing a drug candidate for safety and efficacy can be as simple as applying it to each well of a tissue culture plate populated with cells from many people and watching how the cells respond. If they die or stop functioning, the candidate probably doesn't warrant further investigation. If diseased cells begin to function better, researchers can explore how or why — all without involving laboratory animals or asking human participants to roll up their sleeves or gulp a pill.

Or flip that previous scenario on its head. Test three, or five, or 10 different available drugs on a tray of cells all from one individual and see which is most effective. Now you have an idea as to which treatment might work best for that individual.

The concept can even be extended beyond cells in a dish to include organoids — micro versions of brains or intestines or pancreas made up of several cell types in a tissue.

"We are not going to eliminate the need for safety testing in animals or for clinical trials in people," Wu said. "But if we first test a potential treatment on trays of cells or organoids from hundreds of people in a laboratory, the screening process can be much simpler and more efficient. We can identify subgroups of people who might benefit more than others and quickly weed out drugs with unacceptable safety risks on human cells."

Heart cells from skin cells

CLINICAL TRIALS IN a dish became feasible only when stem cell pioneer Shinya Yamanaka, MD, PhD, showed in 2007 that it's possible to take a specialized cell like a skin cell and wind it backward to its infancy. Unlike mature skin cells, locked into their dermatological destiny, the newly infant cells have potential to become nearly any cell in the body. Scientists deem these cells pluripotent, from the Latin "pluri," meaning many, and "potent," meaning power.

A decade ago, Wu and his colleagues showed that lab-generated cardiomyocytes mirror the genetic profile of their donors' heart muscle. This indicated the cells could be a good proxy on which to test medications for genetic disorders.

It also meant these cells could be used to screen drugs of all sorts for likely impact on the heart. This is important because drugs that fail clinical trials or are pulled from the market after approval are often removed due to cardiac complications — regardless of what condition the drug was intended to treat.

"Now, instead of testing drugs on me as the guinea pig, we can test these drugs on my surrogates — the beating heart cells that are genetically identical to me," Wu said. A tray in which every well contains Wu's iPS-derived heart cells can be easily treated with therapeutic doses of several candidate drugs and assessed for a favorable response.

It's more than a hypothetical scenario. Since 2017, Wu has gotten up close and personal with his own iPS cells — characterizing their responses to various medications. The cellular introspection is driven in part by curiosity, but it's also a testament to what Wu sees as the power of cell-based testing. "Our postdocs have taken my iPS cells and differentiated them into my brain cells, heart cells, endothelial cells and liver cells," Wu said. "I'm asking them to test some of the medications that I might need to take in the future."

Since Yamanaka's 2007 discovery, scientists have shown that iPS cells can also be induced to become many types of cells — meaning that the clinical trial in the dish can be used for many conditions. After you make a genetically identical copy of a person's brain, heart, liver or other type of cell, you only need a way to assess a drug's effect on that cell type.

Clinical trials in a dish can also be used to screen and validate the increasing numbers of drug candidates designed with the help of artificial intelligence. Some of these computational approaches generate dozens or even hundreds of molecules for testing.

"We can now essentially tell a computer program, 'This is the structure of the protein we're targeting; design some molecules that can bind to it,'" Wu said. "And then we can use these clinical trials in a dish to screen these molecules on cells from hundreds of genetically diverse people at one time."

Artificial intelligence-enabled strategies like those Angela Zhang, an MD-PhD student in Wu's lab, is using could facilitate and standardize the use of iPS-derived cells in small labs as well as large pharmaceutical companies. Zhang, who began working in Wu's laboratory as an undergraduate, is using machine- and deep-learning techniques to predict which iPS cell colonies are likely to successfully become functional cardiomyocytes that can be used for screening.

"There's no better place than Stanford to apply machine learning to standardize and scale up these clinical trials in a dish," Zhang said. "And we can do it with cells from diverse backgrounds."

Diversity is a priority for Wu and his colleagues.

"Most of the drugs approved by the FDA in the past were developed and tested primarily on white people, especially men," Wu said. "But we have now the world's largest iPS cell biobank — with cells from more than 2,000 people from a variety of ethnic backgrounds. Some are healthy, some have heart diseases, and others have rare or orphan diseases that are difficult to study. I don't think any other university or single lab has access to a resource of this magnitude."

That biobank is a big part of the reason Stanford Medicine leads the effort to explore the value of clinical trials in a dish. It exists in part due to funding from the National Institutes of Health and the California Institute of Regenerative Medicine and because of people like Wu, who straddle the medical world and the research lab — merging access to patient samples with the lab space and expertise needed to generate such a resource.

"I am a physician scientist," Wu said. "If I only studied mice, my MD might as well stand for mouse doctor. My goal before I retire is to figure out ways to improve our current system of drug development. Right now, time and money are the two biggest obstacles. If I can help decrease the time it takes and the amount of money it costs to get successful treatments to patients, then I will feel that I have made some type of meaningful contribution to biomedical science during my career." **SM** — *Contact Krista Conger at kristac@stanford.edu*

PLUS

An act of love

CONTINUED FROM PAGE 53

just one example of her cultural and religious values impacting her experiences with American medical care. While other Chinese families may react to American health care differently than my mom did, I can see her actions broadly reflected in common cultural conflicts. The Chinese characters in the film *The Farewell* grappled with disclosing a terminal cancer diagnosis to the family matriarch, fearing that the knowledge would irrevocably devastate her and hasten death. For years, my mom refused to acknowledge that the medications my dad took were chemotherapy because she didn't want to think of his brain tumor as malignant, fearing that naming it as such would transform it into cancer. Before every surgery, my mom prayed at the local temple, lighting incense sticks both to protect my dad and to guide his surgeons' hands. She would pull doctors out of his room whenever they discussed disappointing imaging results and the gravity of his terminal condition, lest my dad hear his prognosis and despair.

As a Chinese-American doctor and a daughter with immigrant parents, I understood firsthand how much culture matters in health care. Culture pervades not only customs surrounding death, as I've seen with my dad's last month in the hospital, but also health care delivery, access and communication. The patients and families I encounter at Stanford Hospital arrive with rich and deep experiences and values that could be understood only through patience, empathy and listening. My friend and colleague told me about her interaction with one Chinese family. They couldn't agree on a time to discuss goals of care for their loved one until the conversation was moved to a conference room other than Room 444, for the number 4, in Chinese, sounds like and represents "death." How could we have known that was the problem, unless we thought to ask what the family needed?

Indeed, in the end, it wasn't about my dad's nasogastric tube. It was just a symbol my mom clung to while she faced the devastation of the inevitable. A conduit for his peaceful passing as she struggled to let my father go.

My father had a lovely funeral on a warm, sunny day. My mom chose his final resting place; the gravesite had an impressive vista, overlooking several gently sloping hills with the Los Angeles skyline in the distance.

As Buddhist nuns chanted blessings, we prayed for him before a table laden with flowers and dozens of food platters — rice cakes, chow mein, lotus buns and more. One by one, we ceremoniously raised each dish in the direction of his casket. They were offerings so that his soul would be nourished in heaven. I peeked at Mom, and she had a peaceful expression on her face, the first one I'd seen in months.

Eat well, Dad. We love you. **SM**

Adela Wu, MD, is a neurosurgery resident at Stanford Medicine. Contact her at adelawu@stanford.edu.

Editor:
ROSANNE SPECTOR
Senior Associate Editor:
PATRICIA HANNON
Chief Communications Officer:
CECILIA ARRADAZA
Senior Director, Content Strategy:
ANNIC JOBIN
Art/Design Direction:
DAVID ARMARIO DESIGN
Contributing Editor:
PAUL COSTELLO
Writers:
HANAE ARMITAGE
JONATHAN H. CHEN
KRISTA CONGER
MARK CONLEY
ERIN DIGITALE
BRUCE GOLDMAN
JOHN SANFORD
SARAH C.P. WILLIAMS
ADELA WU
ANNA MARIE YANNY

Video Journalist:
MARK HANLON
Copy Editor:
MANDY ERICKSON
Circulation Manager:
ALISON PETERSON

Stanford Medicine is published by the Stanford Medicine Office of Communications as part of an ongoing program of public information and education.

Stanford MEDICINE

 Requests for subscriptions, address changes or permission to copy or reprint and other correspondence should be emailed to medmag@stanford.edu or addressed to *Stanford Medicine* magazine, Stanford Medicine Office of Communications, 1520 Page Mill Road, MC 5471, Palo Alto, CA 94304
We can be reached by phone at 650-723-6911.

To read the online version of *Stanford Medicine* and to get more news about Stanford Medicine, visit stanmed.stanford.edu. For information from the Stanford Medicine Alumni Association, visit med.stanford.edu/alumni/.

IT WORKED

A PREGNANT WOMAN RISKED A TREATMENT HERSELF, HOPING TO HEAL HER DEVELOPING BABY, WHO WOULD BE BORN WITH CYSTIC FIBROSIS

Like many toddlers, 1-year-old Nora Zalinski's favorite activities are very … well, active. "She loves jumping off of the furniture," her mom, Michaela Zalinski, said with a chuckle. Nora also tries to imitate anything her energetic big brothers, Jaxon, 5, and Elijah, 3, are doing, Zalinski said.

All three children were born with cystic fibrosis, a genetic disease that affects the lungs and digestive system. But in important ways, Nora's medical trajectory has not imitated her brothers'.

Both boys developed severe intestinal blockages in utero, requiring surgery as newborns to remove large portions of injured bowel. When a prenatal ultrasound scan showed a similar blockage emerging in Nora, Zalinski and her husband, Jacob Zalinski, were worried.

But an article they read about a woman with CF who had taken a new treatment for the disease while pregnant gave them hope; the treatment, a three-drug combination called Trikafta, reduced disease signs in her baby at birth. Michaela Zalinski didn't have CF, but might Trikafta help her developing baby too? She hoped so and asked her medical team at Lucile Packard Children's Hospital Stanford if she could try the drug combo.

Michaela Zalinski, here with her three children, took medication while pregnant that fixed a bowel problem arising for Nora.

Approved in 2019, Trikafta is part of a wave of new CF treatments that precisely target the gene mutations underlying the disease. As far as her medical team could tell, no one had treated a mom like Zalinski who carried the CF gene but didn't have the disease and was pregnant with an affected fetus.

"It was a high-risk situation," said Carlos Milla, MD, professor of pediatrics and director of the Cystic Fibrosis Center at Stanford. The boys' severe bowel blockages, the ultrasound showing the same problem developing a third time and the option to try something new were all compelling, he said, adding, "In the past, we could intervene only after the baby was born."

"Our primary concern was that we could injure this mother who otherwise would not be taking Trikafta," said Susan Hintz, MD, professor of pediatrics and medical director of the Fetal and Pregnancy Health Program at Stanford Medicine Children's Health. After weighing possible side effects of the drugs, the doctors went ahead. "It was a heroic effort on the team's part," said Natali Aziz, MD, Michaela's high-risk obstetrician at Stanford Medicine Children's Health.

Zalinski didn't start taking Trikafta until more than halfway through pregnancy, which worried the family — as the previous mothers with CF had taken it from the outset. "We had no clue if it was actually going to do anything for us," she said. "We were just hoping."

On further prenatal scans, Nora's blockage seemed to disappear, a better result than the physicians expected. They warned that this might mean Nora's bowel had perforated; they wouldn't know until she was born.

Right after her birth on March 12, 2022, the medical team performed an X-ray with a contrast agent to see Nora's digestive system. Everything was healthy. Unlike her brothers, she did not need surgery.

"We were so relieved," Zalinski said. Both boys required months of hospitalization after birth; Jacob and Michaela, thrilled when 4-day-old Nora was released from the hospital, giddily learned to adjust a car seat for a newborn.

"Our team is so encouraged to know there is a possibility to shift patients like Nora to a much easier road," Hintz said. "For all of us, that is the most exciting part of this work." BY ERIN DIGITALE

COURTESY OF THE ZALINSKI FAMILY

S T A N F O R D
M E D I C I N E

Stanford Medicine
Office of Communications
1520 Page Mill Road, MC 5471
Palo Alto, CA 94304

Where is 'I'?

THE PHYSICAL SENSE OF SELF HAS BEEN LOCALIZED

Ever wonder where in your brain that interesting character called "I" lives? Stanford Medicine physician-scientist Josef Parvizi, MD, PhD, has ways of peeking into people's heads and finding out what makes us tick. Parvizi's most recent sighting, described in a June 8 paper in *Neuron*, unveils the surprising role of a small structure sandwiched between the brain's two hemispheres.

The structure, called the anterior precuneus, is the hub of a network of brain regions whose activity rises and falls in coordination with one another's, indicating teamwork. Parvizi and his colleagues discovered that this archipelago of collaborating brain regions spearheaded by the anterior precuneus is key for integrating information about your location, motion, sensations, and muscle and joint positions to form a mental map of your sense of bodily, or physical, self.

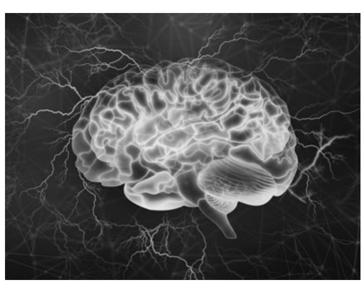

To explain the network's significance, Parvizi invoked that odd couple, "I" and "me."

"For every action we take, even during dreams," he said, "there's always an agent behind it: We call that agent 'I.'" It's the physical/bodily sense of self. "'Me' is everything we have stored in our memories about the 'I'," said Parvizi, a professor of neurology. It's the narrative self that actively or passively thinks about things like memories, habits, emotions and plans.

Distinct assemblages of brain structures govern the two systems, which continuously interact with each other. "Me" dwells in a well-studied network of neurons in the brain called the default mode network. There's no official name yet for the "I" network.

Parvizi's team delineated the anterior precuneus's role through studies of epilepsy patients who'd had electrodes implanted in their brains to help locate their seizures' origin. The team found that stimulating the network led by that brain structure wreaked havoc on the patients' ability to place themselves in space.

In contrast, electrically stimulating the default mode "me" network does nothing at all to one's physical sense of self, Parvizi said. This shows these are two entirely separate, although interacting, networks, he said.

"Your sense of physical or bodily self represents your organism in the immediate here and now, with a particular point of view that is yours alone, your first-person perspective on the world around you. Nobody shares it," Parvizi said. "You may not be conscious of your point of view. But you will be if I disrupt the network that generates it. Your place in the world around you will suddenly seem unreal." — BRUCE GOLDMAN

SIARHEI