THE JOHNS HOPKINS WHITING SCHOOL OF ENGINEERING MAGAZINE

JHUENGINEERING

SPRING 2024

Algorithms for a Fairer World

Addressing bias in AI applications.

In AI We Trust?

Our experts share insights.

Delivering on Precision Health

Harnessing AI advances to improve health care.

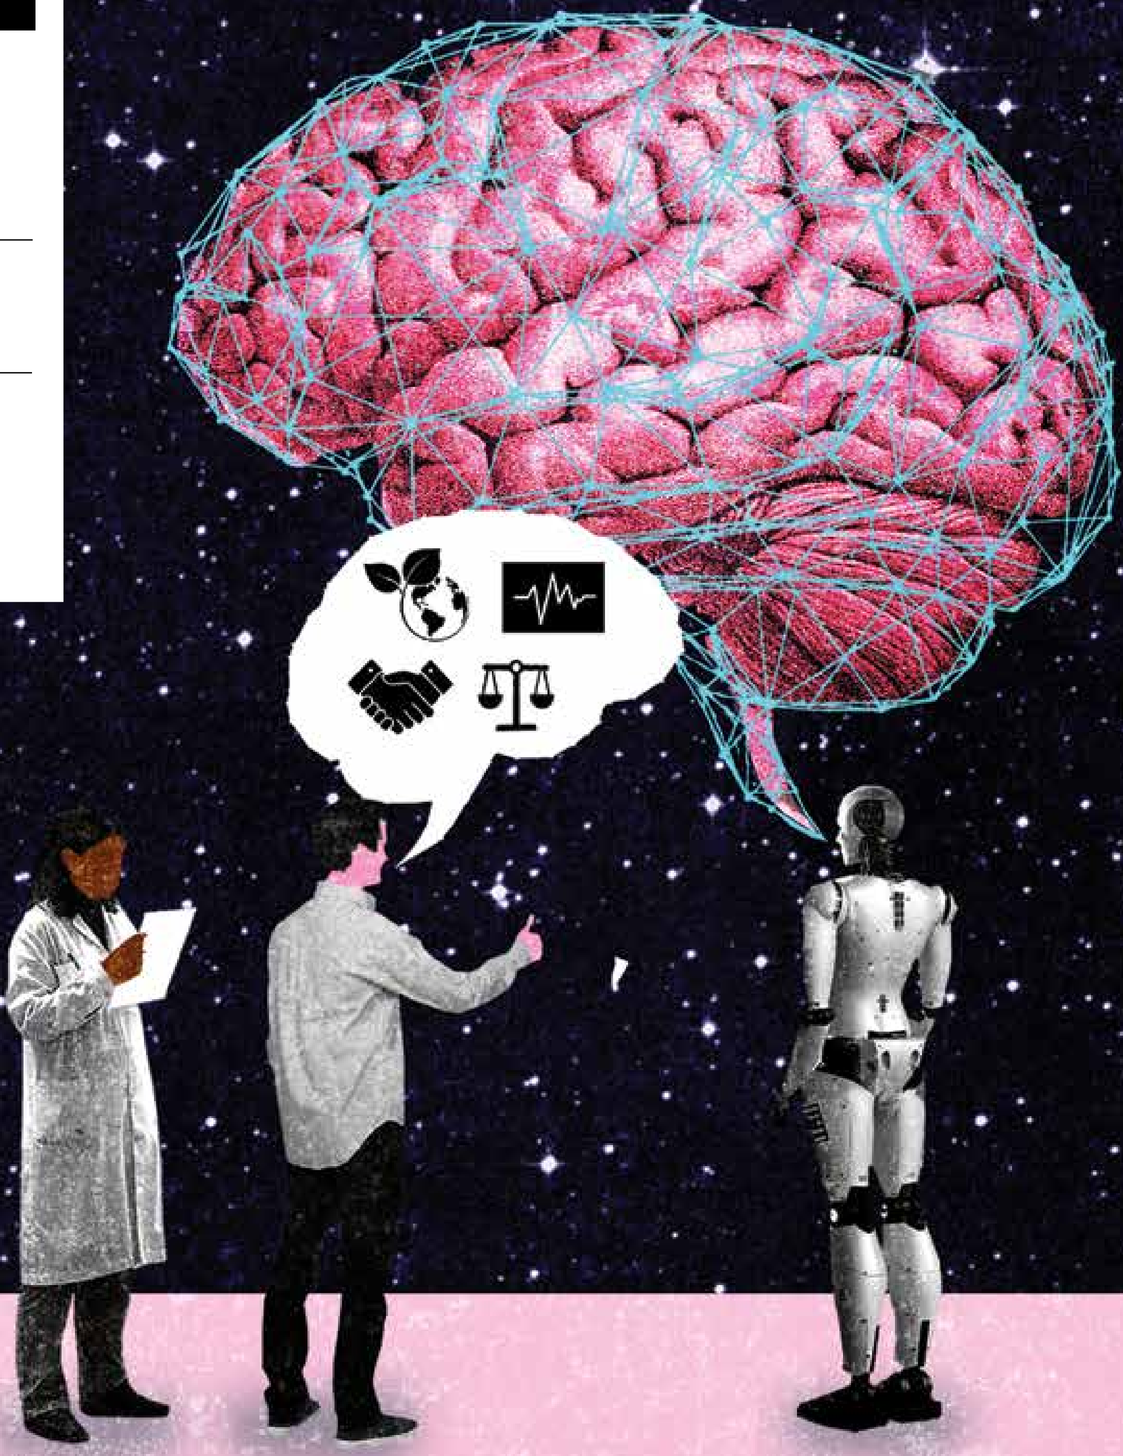

Special Issue:

Artificial intelligence advances are rocking our world. What's the promise—and potential peril?

FROM THE DEAN

Dear WSE community,

WE ARE COMMITTED TO CREATING A COMMUNITY WHERE EVERY MEMBER CAN USE THEIR TALENTS AND hard work to achieve their goals—however ambitious they may be.

Our school is not our programs or the buildings in which we offer them; rather, it is the people who support each other's aspirations and our mission. Our goal is to inspire ourselves and others to think deeply and challenge our beliefs, and this can only happen in an environment where every member of our community knows they are valued and their voices are heard.

We recently surveyed faculty, students, and staff members to better understand their perceptions of our success in creating such an inclusive environment and to provide us with information to guide our ongoing efforts. While the overall results suggest that the members of our community feel satisfied and supported by the school, the results also highlight areas where we need to improve to enhance the experiences of all.

We also are making great strides in achieving the vision we defined two years ago to achieve preeminence in the areas of artificial intelligence, data science, bioengineering, and energy and the environment and are actively expanding the size of our faculty and building infrastructure in these key areas. Our goal is not simply to grow bigger, but to grow in a way that increases our impact on the world and our leadership in engineering research and education.

To achieve this, we think carefully about how we operate and make adjustments as needed. As a result of the Second Commission on Undergraduate Education's findings, we are implementing changes in our approach to undergraduate education. To better meet the needs of our community, we have established a Whiting School Staff Council, Office of Human Resources, and Office of Diversity and Inclusion, and are exploring a new shared governance structure.

As the school continues to evolve and grow, I am confident that we will continue to build structures and policies that will allow us to adapt to changing times and needs and ensure our success. I look forward to sharing our progress with you.

Best wishes,

ED SCHLESINGER
Benjamin T. Rome Dean

CONTRIBUTORS

MIKE FIELD

ALGORITHMS FOR A FAIRER WORLD (P. 12)

Field, a freelance journalist whose work has appeared both nationally and internationally, holds an MA in Drama Studies from Johns Hopkins. His article, "Asthma, The Breathtaking Disease" (*Hopkins Public Health Magazine*), won a gold medal for feature articles from the Council for Advancement and Support of Education.

PATRIC SANDRI

SPECIAL REPORT: ADVANCES IN AI (P. 12, 20, 28)

Sandri, who holds an MA from the Royal College of Art in London, lives and works in Zurich, Switzerland. His conceptual illustration—incorporating collage, digital drawing, and painting—has been published in *TIME* magazine, *The Wall Street Journal, Boston Globe,* and *Advertising Age,* among others. His paintings have been exhibited at numerous galleries in Basel, Berlin, Vienna, and Madrid.

ABOUT THIS MAGAZINE

Editor
Abby Lattes

Managing Editor
Lisa Ercolano

Consulting Editor
Sue De Pasquale

Art Direction and Design
Cut Once, Inc. (cutoncedesign.com)

Associate Dean for External Relations
Megan Howie

Director of Communications
Abby Lattes

Director of Constituent Engagement
Kim Sheehan Dolan A&S '13, MS '18

Contributing Writers:
Conner Allen, Nikolas Charles, Jonathan Deutschman, Wick Eisenberg, Lisa Ercolano, Mike Field, Adam Hadhazy, Dino Lencioni, Robin Smith, and Abigail Sussman

Contributing Photographers and Illustrators:
Will Kirk A&S '99/Homewood Photo, and Patrick Sandri

JHU Engineering magazine is published twice annually by the Whiting School of Engineering Office of Marketing and Communications.

We encourage your comments and feedback.
Please contact us at:
WSE Communications
info_wse@communication.jh.edu

To update your mailing address or order additional copies, email engineering@jhu.edu or call 410-516-8723.

JOIN THE DISCUSSION ONLINE:

facebook.com/jhuengineering
twitter.com/HopkinsEngineer

DEAN SCHLESINGER PHOTO: WILL KIRK/HOMEWOODPHOTO.JHU.EDU.

CONTENTS

Cover illustration by Patric Sandri.

AT WSE

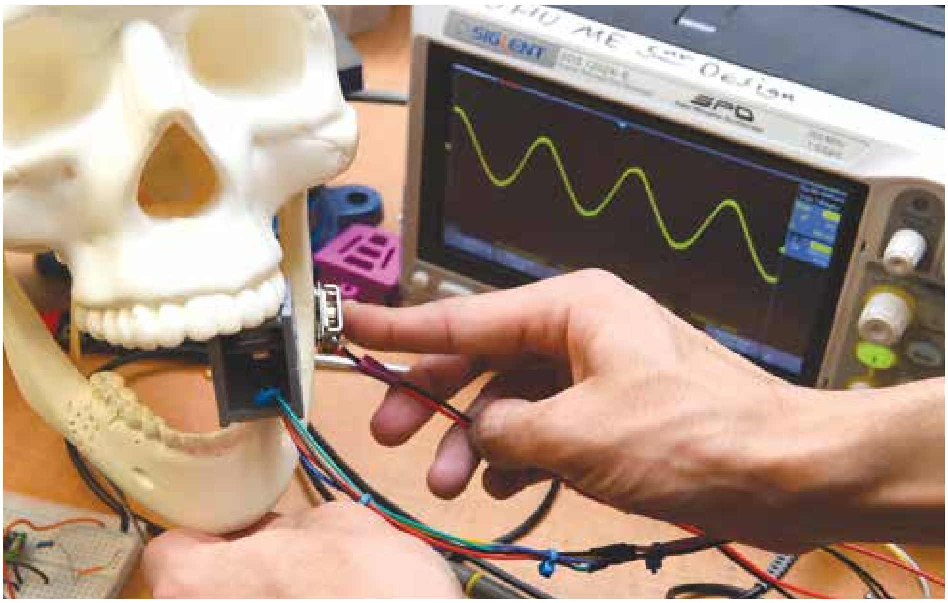

Designing Solutions

MORE THAN 600 WHITING SCHOOL OF ENGINEERING students participated in this year's Design Day, the school's annual celebration of innovation and creativity. Held on May 1, the event showcased students' ability to apply theoretical knowledge gained in classrooms and laboratories to real-world problems.

This year's projects ranged from an acoustic device designed to attract (and track) oysters in the Chesapeake Bay to a migraine-fighting electrostimulation headband to an AI system to detect flaws in fabric used to make uniforms for the U.S. military—and much more, including:

DAMPENING DENTAL DRILL DIN

A team of mechanical engineering students worked with a Florida dentist to develop a device that dampens the buzzing and sometimes frightening sound of the drill used during procedures such as root canals (above left). Using an anatomical skull and their own mouths as sound conductors, the students created a bite plate that uses the same technology as noise-canceling headphones. The device picks up vibrations and produces an "anti-noise" wave; an attached controller allows the patients to adjust during treatment.

FLEXING MUSSELS FOR WATER QUALITY

A team of environmental health and engineering students worked with Baltimore City to develop a plan to install mussel beds in Back River (top right) and oyster beds in the Patapsco River. The shellfish, which will be nursed in cages before being transferred to reefs downstream, will serve as natural filters, reducing nutrient pollution in the water.

ANALYZING EVERYDAY MOTION TO TRACK PARKINSON'S PROGRESS

Biomedical engineering students have developed a signal-processing algorithm to help clinicians understand how well medication is working for patients with Parkinson's, a neurodegenerative disease that affects an estimated 9 million patients worldwide (above bottom). Using small motion-tracking devices called inertial measurement units (IMUs), the team will record patients' movements as they complete everyday tasks, such as brushing their teeth or writing with a pen.

THIS PAGE: WILL KIRK/HOMEWOODPHOTO.JHU.EDU.

AWARDS AND HONORS

NATIONAL ACADEMY OF INVENTORS FELLOW

HAI-QUAN MAO, professor of materials science and engineering and director of the Institute for NanoBioTechnology, was elected a Fellow of the National Academy of Inventors. A pioneer in the field of regenerative medicine and therapy delivery, Mao creates nanofiber scaffolds to regenerate tissues such as the liver, soft tissues, and blood vessels.

EARLY CAREER AWARD

SOLEDAD VILLAR, assistant professor of applied mathematics and statistics, received a National Science Foundation Early CAREER Award, which recognizes early-stage scholars with high levels of promise and excellence, for her project "Symmetries and Classical Physics in Machine Learning for Science and Engineering."

ROSALIND FRANKLIN YOUNG INVESTIGATOR AWARD

YUTING LUO, assistant professor in materials science and engineering and a member of the Ralph O'Connor Sustainable Energy Institute, received the Argonne National Laboratory's 2024 Rosalind Franklin Young Investigator Award for her materials and electrochemistry research.

NEW LEADERSHIP FOR DOCTOR OF ENGINEERING PROGRAM

ASHUTOSH DUTTA, the chief 5G strategist at the Johns Hopkins Applied Physics Laboratory and an IEEE Fellow, was named director of the Whiting School's Doctor of Engineering program.

Established in 2018, the Doctor of Engineering program, which is designed for working professionals, includes research collaboration between a student's employer and the Whiting School.

Dutta is focusing on growing the program by increasing the number of APL staff members and Johns Hopkins faculty serving as program advisers, establishing new partnerships with industry and government to increase corporate engagement and research collaboration, and growing annual enrollment by enhancing the resources available to prospective applicants.

Dutta previously served as chair of the Engineering for Professionals electrical and computer engineering program and is a member of JHU's Institute for Assured Autonomy.

ONLINE GRADUATE PROGRAM AMONG NATION'S BEST

Johns Hopkins Engineering's online graduate degree programs continue to be recognized by *U.S. News & World Report* as among the nation's best.

In the publication's 2024 rankings, announced in February, multiple EP programs ranked No. 1, including its online computer information technology offerings, which tied for the top spot. Rankings in this category considered data from five EP fields: computer science, artificial intelligence, cybersecurity, information systems engineering, and data science.

Also tied for the No. 1 spot are EP's engineering management and electrical and computer engineering programs.

EP's online programs overall now tie for No. 6 among the Best Online Master's in Engineering Programs. This year's specialty rankings include EP's industrial/systems engineering program at No. 8; mechanical engineering at No. 7; and civil engineering, which tied for No. 7.

— LE

Shoring Up Wind Energy

RESEARCHERS FROM JOHNS HOPKINS' RALPH O'CONNOR SUSTAINABLE ENERGY INSTITUTE (ROSEI) AND MORGAN STATE UNIVERSITY ARE teaming up with the State of Maryland's Maryland Energy Administration (MEA) to establish the state's presence in a new national center supporting offshore wind energy.

The effort—established through the U.S. Department of Energy's Academic Center for Reliability and Resilience of Offshore Wind (ARROW)—aims to expand the local offshore wind industry and provide greater opportunities for Maryland residents and businesses to participate in the growing clean energy economy.

"By being a member of this new national center, MEA, Hopkins, and Morgan State are announcing that we are in it for the long haul to make the energy transition work with offshore wind in Maryland," says Ben Schafer, director of ROSEI and a professor of civil and systems engineering at the Whiting School of Engineering.

— WICK EISENBERG

WATERMAN AWARD


Muyinatu "Bisi" Bell

MUYINATU "BISI" BELL, John C. Malone Associate Professor in electrical and computer engineering, received the National Science Foundation's 2024 Alan T. Waterman Award, considered the nation's highest honor for early-career scientists and engineers, for her work "harnessing light and sound to improve biomedical imaging systems."

As director of the Photoacoustic and Ultrasonic Systems Engineering (PULSE) Lab, Bell is advancing health care interventions and diagnosis through techniques that are transforming ultrasound and photoacoustic imaging. Among her innovations are an algorithm that can enable accurate ultrasound images, regardless of a patient's skin tone or BMI (see p. 30), and an AI-driven diagnostic tool that can identify COVID-19 in lung ultrasounds. Bell has joint appointments in the departments of Biomedical Engineering and Computer Science.

Turning Manure Into Fuel

CHEMICALS—FROM ANTIBIOTICS USED TO KEEP LIVESTOCK HEALTHY to pesticides that shield crops from insects—play an important role in modern agriculture. However, many of these substances accumulate in cow manure used as fertilizer, where they contaminate crops, leach into groundwater, and pollute waterways.

A study led by Johns Hopkins environmental engineers suggests that anaerobic digestion—using microbes to break down organic matter in the absence of oxygen—could remove these potentially harmful chemicals from cattle manure before they pollute nearby soil and water. What's more, this process can also produce sustainable energy because the organic matter in the manure breaks down to form methane, a bio-gas.

"The increasing use of antibiotics and pesticides in agriculture present significant risks to human health and the environment, making it crucial to come up with measures to mitigate their negative impact," says team member **Carsten Prasse**, an assistant professor of environmental health and engineering. "Our work spotlights the potential of anaerobic digestion as one solution." The team's results appeared in *Science of the Total Environment*.

Anaerobic digestion is commonly used to treat sewage sludge in the U.S. It is also gaining popularity in other countries as a method for processing livestock manure. The goal of these approaches is to reduce greenhouse gas emissions from the offal, typically stored in large open lagoons. Previous research has focused primarily on the impact of veterinary antibiotics on microorganisms involved in anaerobic digestion.

Prasse's team had a different aim: determining whether anaerobic digestion could be used to remove these chemicals as well as potentially produce bio-gas. Furthermore, instead of focusing solely on whether anaerobic digestion removed the chemicals, the researchers investigated the substances created during the anaerobic removal process.

"By investigating the pathways by which the microbes broke down the chemicals, we gained insight into what new substances were produced," Prasse says. "This sheds light on which microbes facilitated digestion and whether they succeeded in fully breaking down the 'parent' contaminants or merely transformed them into other, new, and potentially concerning compounds."

In their study, the researchers investigated the fate of 20 common cattle-waste substances in small laboratory containers filled with manure as they underwent anaerobic digestion. They found that 11 of those substances broke down, creating 47 new ones (called "transformation products") in the process.

The team also made another important discovery: Removing chlorine from organic compounds results in less persistent and often less toxic transformation products.

— LISA ERCOLANO

OPPOSITE PAGE: ILLUSTRATION IS A STILL FROM AN ANIMATION BY SURYANSH PRAKHAR.

TELL US WHAT YOU THINK!

Continue the conversation online, using **#JHUEngineering** on Facebook and Twitter!

Reimagining a Quieter Drone

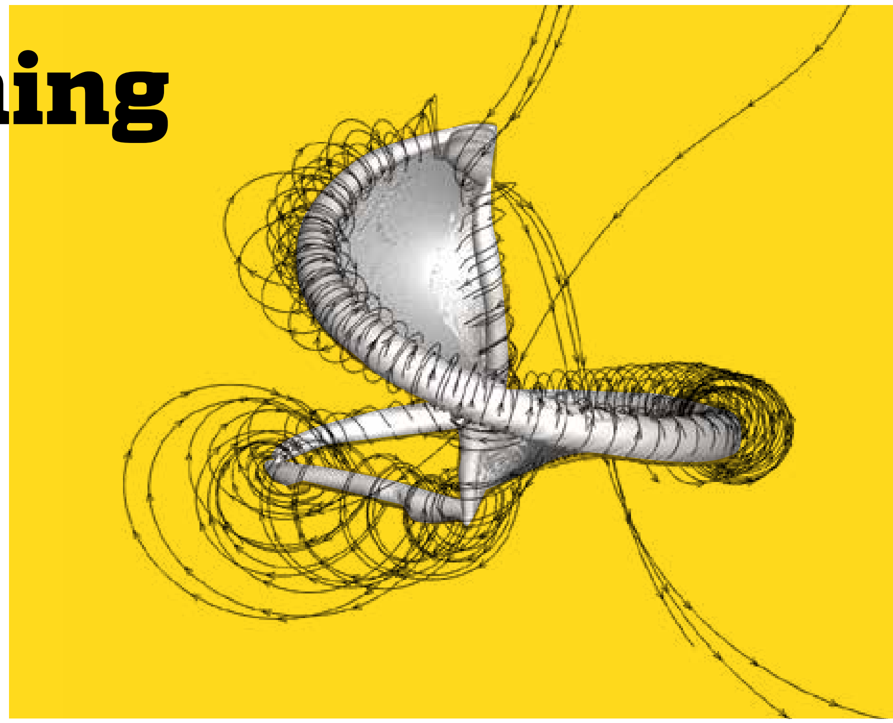

THE FAA RECEIVES THOUSANDS OF COMPLAINTS EVERY YEAR ABOUT THE LOUD BUZZING PRODUCED by drone propellers slicing through the air. Experts believe noise pollution caused by these unmanned aerial vehicles will only worsen as they are increasingly used for package delivery, emergency response, and more.

A research team led by mechanical engineering professor **Rajat Mittal** believes a device invented by Leonardo da Vinci more than 500 years ago may hold the key to quieter drones. "da Vinci's visionary aerial screw—a sort of precursor to the modern helicopter—inspired our investigation," says Mittal, whose mechanical engineering team included **Jung-Hee Seo**, an associate research professor, and doctoral candidate **Suryansh Prakhar**. The researchers published their findings in *Bulletin of the American Physical Society*.

The team was aware that other research groups had explored loop-shaped propellers that were less noisy than traditional propellers that have flat, thin blades and angled edges. The characteristic buzzing sound produced by traditional propellers is the result of "tip vortices"—small swirling tornadoes of air that whoosh and intersect with the flat, angled blades. Loop propellers spread those vortices around, muting the sound.

Mittal's team thought that da Vinci's design, with its screw-like shape and single blade, might be quieter, too. To start, the researchers created a model, building on the work of aerospace engineering students at the University of Maryland, who had analyzed the aerial screw's design, including its radius, curve, pitch, shape, and number of loops.

"We constructed a 3D model of the da Vinci aerial screw's looped shape, and then used our simulation software, called ViCar3D, to simulate the flow of air around the rotor as the drone was hovering in place. The software then predicted the speed of airflow around the propeller and pressure patterns," says Prakhar.

The pressure generated on the rotating screw's surface would turn into sound, so the team calculated the noise produced five meters from the rotor. Then they simulated the loop propeller in the same circumstances.

The da Vinci propeller was, in fact, a bit noisier than the loop propeller at any given rotation speed, but the aerial screw also produced more lift. Knowing the amount of lift needed to remain fixed for common drone assignments, the team then calculated the noise emitted by the da Vinci and loop propellers when producing the same amount of lift and had their answer.

"The da Vinci propeller produced much less noise for the same given amount of lift being generated," Prakhar says.

— LE

Leonardo da Vinci's aerial screw.

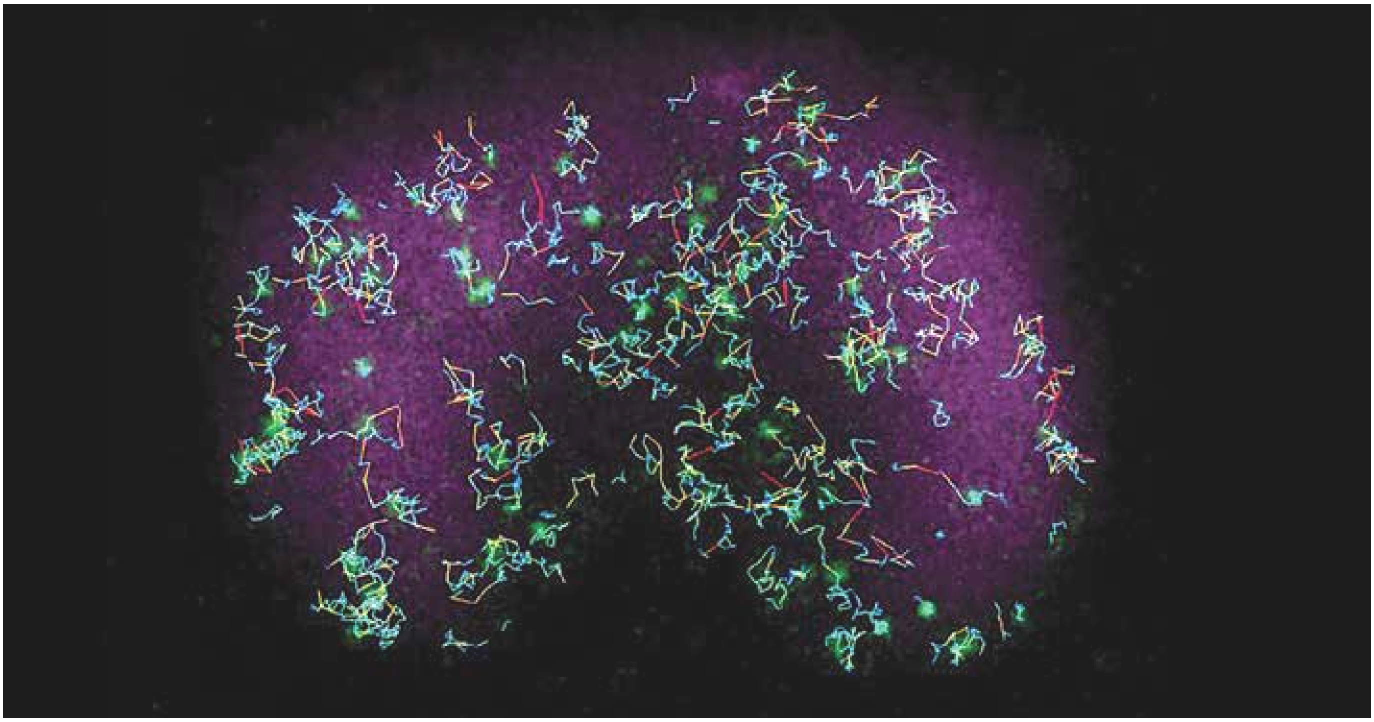

dynamic partitioning-740x494

Bright Ideas in Cellular Dynamics

SCIENTISTS HAVE LONG KNOWN THAT THE CELL MEMBRANE—THE THIN LAYER SURROUNDING EACH cell—plays a crucial role in cell activities and communications. What hasn't been clear, though, is how various proteins that are part of this membrane can dynamically change their location and grouping on the cell surface in response to various signals.

Understanding this mechanism could provide a deeper comprehension of cellular dynamics in general and also pave the way for targeted medical interventions for diseases related to cell signaling and migration.

A team of Johns Hopkins researchers—collaborating with researchers from Osaka University and RIKEN Center for Biosystems Dynamics Research, Japan—discovered a new method for selectively sorting certain proteins into distinct regions of cell membranes, providing insights into how cellular components organize themselves on a large scale during different physiological activities.

The researchers introduced a mechanism called "dynamic partitioning," which operates on the idea that specific proteins can naturally or in response to external signals gather or organize themselves on the cell membrane. This occurs by adjusting their movement in different areas of the membrane.

"This discovery helps us to finally understand how proteins end up in different parts of the cell membrane and dynamically readjust their location as and when required. Our results provide a new lens through which we can understand how signaling events are tightly controlled, as cells communicate, migrate, or divide," says team leader **Pablo A. Iglesias**, the Edward J. Schaefer Professor of Electrical and Computer Engineering at the Whiting School and interim head of that department.

The team started the project after discovering that certain proteins, believed to be uniformly spread across the cell membrane, actually consistently undergo dynamic polarization during various cellular events. "We were intrigued by this surprising result and wanted to dig deeper," says Tatsat Banerjee, PhD '24, lead author of the study, which appeared in *Nature Communications*, and former PhD student in the Whiting School's Department of Chemical and Biomolecular Engineering and the School of Medicine's Department of Cell Biology.

Banerjee and his team then used advanced microscopy techniques to observe and control events inside live *Dictyostelium* amoeba and mammalian immune cells in real time, with detailed resolution at the subcellular level. When they combined computer simulations and real-life experimental results, they found that many proteins in cells move selectively based on the organization of the membrane and how specific proteins sense and interact with its spatial variations.

— DINO LENCIONI

THIS PAGE: ILLUSTRATION COURTESY OF PABLO A. IGLESIAS.

"Though there is a long way to go, the immediate application for our materials is to produce a retinal prosthetic that can be used to produce what we call a 'vision-like' experience for patients suffering from a loss of vision."

— ALEXANDRA PATTERSON

A Retinal Prosthetic to Restore Sight?

RETINITIS PIGMENTOSA GRADUALLY STEALS THE EYESIGHT OF AN ESTIMATED 2 MILLION PEOPLE around the world. Currently, there are no effective treatments to stop or correct the damage caused by this inherited ailment, which kills off photoreceptor cells in the back of the eye.

A groundbreaking material developed by Johns Hopkins researchers could change that. When made into a retinal implant and activated by laser illumination, the biocompatible nanocomposite material shows promise in acoustically stimulating cells in the diseased retina, sending signals to the brain's visual cortex.

Led by **Emad Boctor**, an associate research scientist at the Malone Center for Engineering in Healthcare, and Johns Hopkins University Applied Physics Laboratory's Seth Billings, a computer scientist and electrical engineer in APL's Research and Exploratory Development Department, the team also included researchers at the Wilmer Eye Institute.

"Though there is a long way to go, the immediate application for our materials is to produce a retinal prosthetic that can be used to produce what we call a 'vision-like' experience for patients suffering from a loss of vision," says **Alexandra Patterson**, a graduate student in the Whiting School's Department of Materials Science and Engineering, who worked on the project with Materials Science and Engineering Professor **James Spicer**. "But there are also potential applications for this material in other areas of medicine, as well as in energy production, precision measurements, and imaging."

The soft, flexible material developed by the team contains nanoparticles of palladium, which can convert light into sound waves that activate retinal cells. Using a technique that let them grow the particles directly inside a silicone matrix ensured that the resultant material would be capable of producing strong, stimulating sound waves when exposed to low-power pulses of light.

The team used X-ray diffraction and transmission electron microscopy to confirm the size and arrangement of palladium particles within the silicone matrix and then assessed its optical properties using a spectrophotometer, a tool that measures how a sample absorbs or transmits light.

"We learned that for this application, the materials must be biocompatible, stable, and robust when exposed to the laser—a set of characteristics that is challenging to achieve," says Patterson. "Enhanced performance in one area can mean degraded performance in another."

Currently, other members of the team are testing how the material interacts with retinal cells, while Patterson and Spicer are exploring transforming their novel material into a specific curved shape.

— LE

Turning the Tide Against Superbugs

Calling all superbug slayers! Meet RoboDrop, the miniature terminator dishing out droplet combos to decimate drug-resistant bacteria. Developed by a team led by **Tza-Huei "Jeff" Wang**, the Louis M. Sardella Professor in Mechanical Engineering and member of the Johns Hopkins Institute for NanoBioTechnology, this robotic platform rapidly screens thousands of antibiotic mixtures, identifying those with the most potent punches against stubborn superbugs.

In experiments, RoboDrop discovered a three-drug cocktail that eliminated an *E. coli* strain, even at concentrations where individual antibiotics failed. It precisely prints complex droplet arrangements onto Petri dishes using samples 20,000 times smaller than a teardrop.Customized software also lets it run experiments autonomously.

"We envision this driving the development of new treatment strategies and the optimization of existing antibiotics, turning the tide against superbugs and saving countless lives threatened by antimicrobial resistance worldwide," says collaborator **Fangchi Shao**, a doctoral candidate at the School of Medicine.

— LE

SCAN TO LEARN MORE...

TELL US WHAT YOU THINK!

Continue the conversation online, using **#JHUEngineering** on Facebook and Twitter!

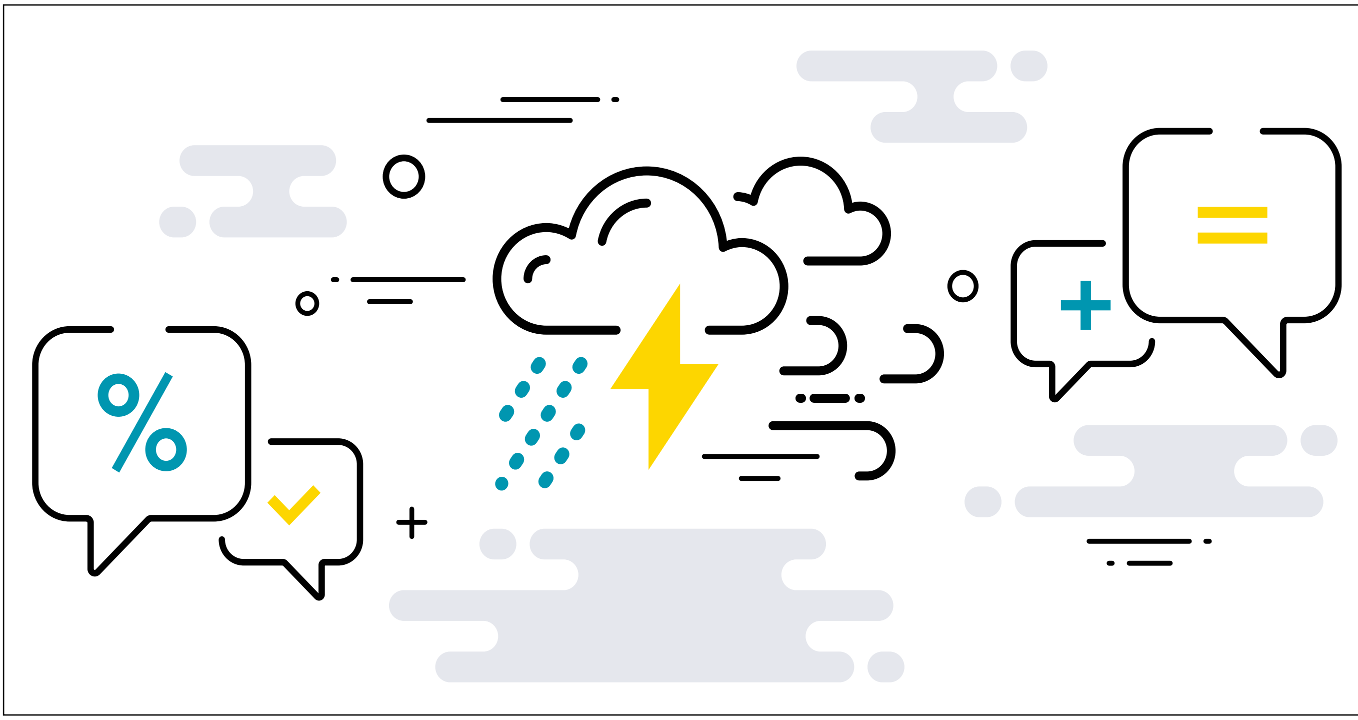

Clearer Predictions for River Flood Damage

RIVERS OVERFLOWING THEIR BANKS HAVE CAUSED PROPERTY DAMAGE AND LOSS OF LIFE IN TEXAS, China, and Kenya in recent months, underscoring the increasing perils of climate change.

To prepare for such floods, governments deploy mathematical models. However, due to time constraints and lack of data, these models sometimes incorporate "off-the-shelf" damage calculations based on previous unrelated floods, which are often inaccurate.

Gonzalo Pita, an expert on natural disaster risk modeling, has developed a reliable and affordable way for governments to estimate expected damage from river floods. This new method provides users with step-by-step instructions and also measures and assigns numerical values to the level of uncertainty in individual flood damage forecasts, giving governments a clearer picture of how reliable their predictions are.

"Accurate predictions are crucial to the safety and well-being of people and property. If a government acts based on inaccurate information, its preparation can be off by orders of magnitude, with very serious results," says Pita, an associate research scientist in the Department of Civil and Systems Engineering and an instructor in the Johns Hopkins Engineering for Professionals' civil engineering program.

The study appears in *The International Journal of Disaster Risk Reduction*. It builds upon work that previously appeared in *The Journal of Hydrology*.

In the new study, Pita first investigated the accuracy of using expert opinion alone to estimate and predict flood damage. He surveyed multiple authorities and simulated thousands of expert surveys in numerous combinations to analyze how the composition of the expert team influences prediction accuracy.

Pita then tackled the issue of "damage functions," a fundamental component of natural disaster risk simulations. "With this method, these functions can be built inexpensively but with a useful level of accuracy that governments can use provisionally until they get better data that will enable them to generate more accurate functions," he says.

The new approach works by helping quantify the uncertainty in experts' predictions by assigning weights to each expert, resulting in a more detailed analysis of the uncertainties involved. The result is a method that Pita expects to be "very useful" for flood modelers and influencing preparedness policy, resulting in "insights that could inform policy directly and indirectly—from enabling smarter zoning laws and budgeting for asset maintenance to designing disaster insurance programs."

— LE

OPPOSITE PAGE: WILL KIRK/HOMEWOODPHOTO.JHU.EDU.

“Nothing is staying put in the ocean. Everything is moving.”

3/26/24 ***USA Today***
James Bellingham, Bloomberg Distinguished Professor of Exploration Robotics, on why rescue efforts were so challenging after the Francis Scott Key Bridge collapsed.

“This is very inspiring for not only my group but for groups around the world to start thinking in this direction when designing technology. Does it serve the wider population?

10/16/23 CNN
Muyinatu “Bisi” Bell, John C. Malone Associate Professor in Electrical and Computer Engineering, on her algorithm that improves ultrasound imaging on patients of all skin tones and body sizes. (Read p. 30 for more.)

“The lead time on air conditioning equipment right now for a home renovation is like 16 months, right? So it’s like you’re telling me they’re going to build a whole bridge in two years? I want it to be true, but I think empirically it doesn’t feel right to me.”

3/29/24 Associated Press
Ben Schafer, Willard and Lillian Hackerman Professor of Civil and Systems Engineering and founding director of the Ralph O’Connor Sustainable Energy Institute, on how long it may take to replace Baltimore’s Francis Scott Key Bridge.

“Freeform use of synthesizing someone’s voice is a slippery slope.”

2/20/24 Scripps News
Anton Dahbura, Computer Science and co-director of the Institute for Assured Autonomy, about efforts to use AI to recreate the voices of people killed by gun violence.

JOHNS HOPKINS
WHITING SCHOOL
of ENGINEERING

JOHNS HOPKINS ENGINEERING
DOCTOR OF ENGINEERING

The One-of-a-Kind *Doctoral Degree* for Working Professionals

EARN YOUR DOCTORATE WHILE YOU WORK and collaborate on a corporate-focused research project with your employer along with expert faculty at Johns Hopkins Engineering.

It's one of the most innovative doctoral degree programs in the field—delivering advanced training for working professionals while opening access to the innovation and research of Johns Hopkins.

Corporate Supported

Mentored Success

Win/Win Results

Learn more about our
DEng Program Today at **Deng@jhu.edu**

engineering.jhu.edu/doctor-of-engineering

ALGORITHMS FOR A Fairer World

BY Mike Field ILLUSTRATIONS BY Patric Sandri

Machine learning technologies hold the potential to revolutionize decision-making. But how can we ensure AI systems are free of bias? Our experts weigh in.

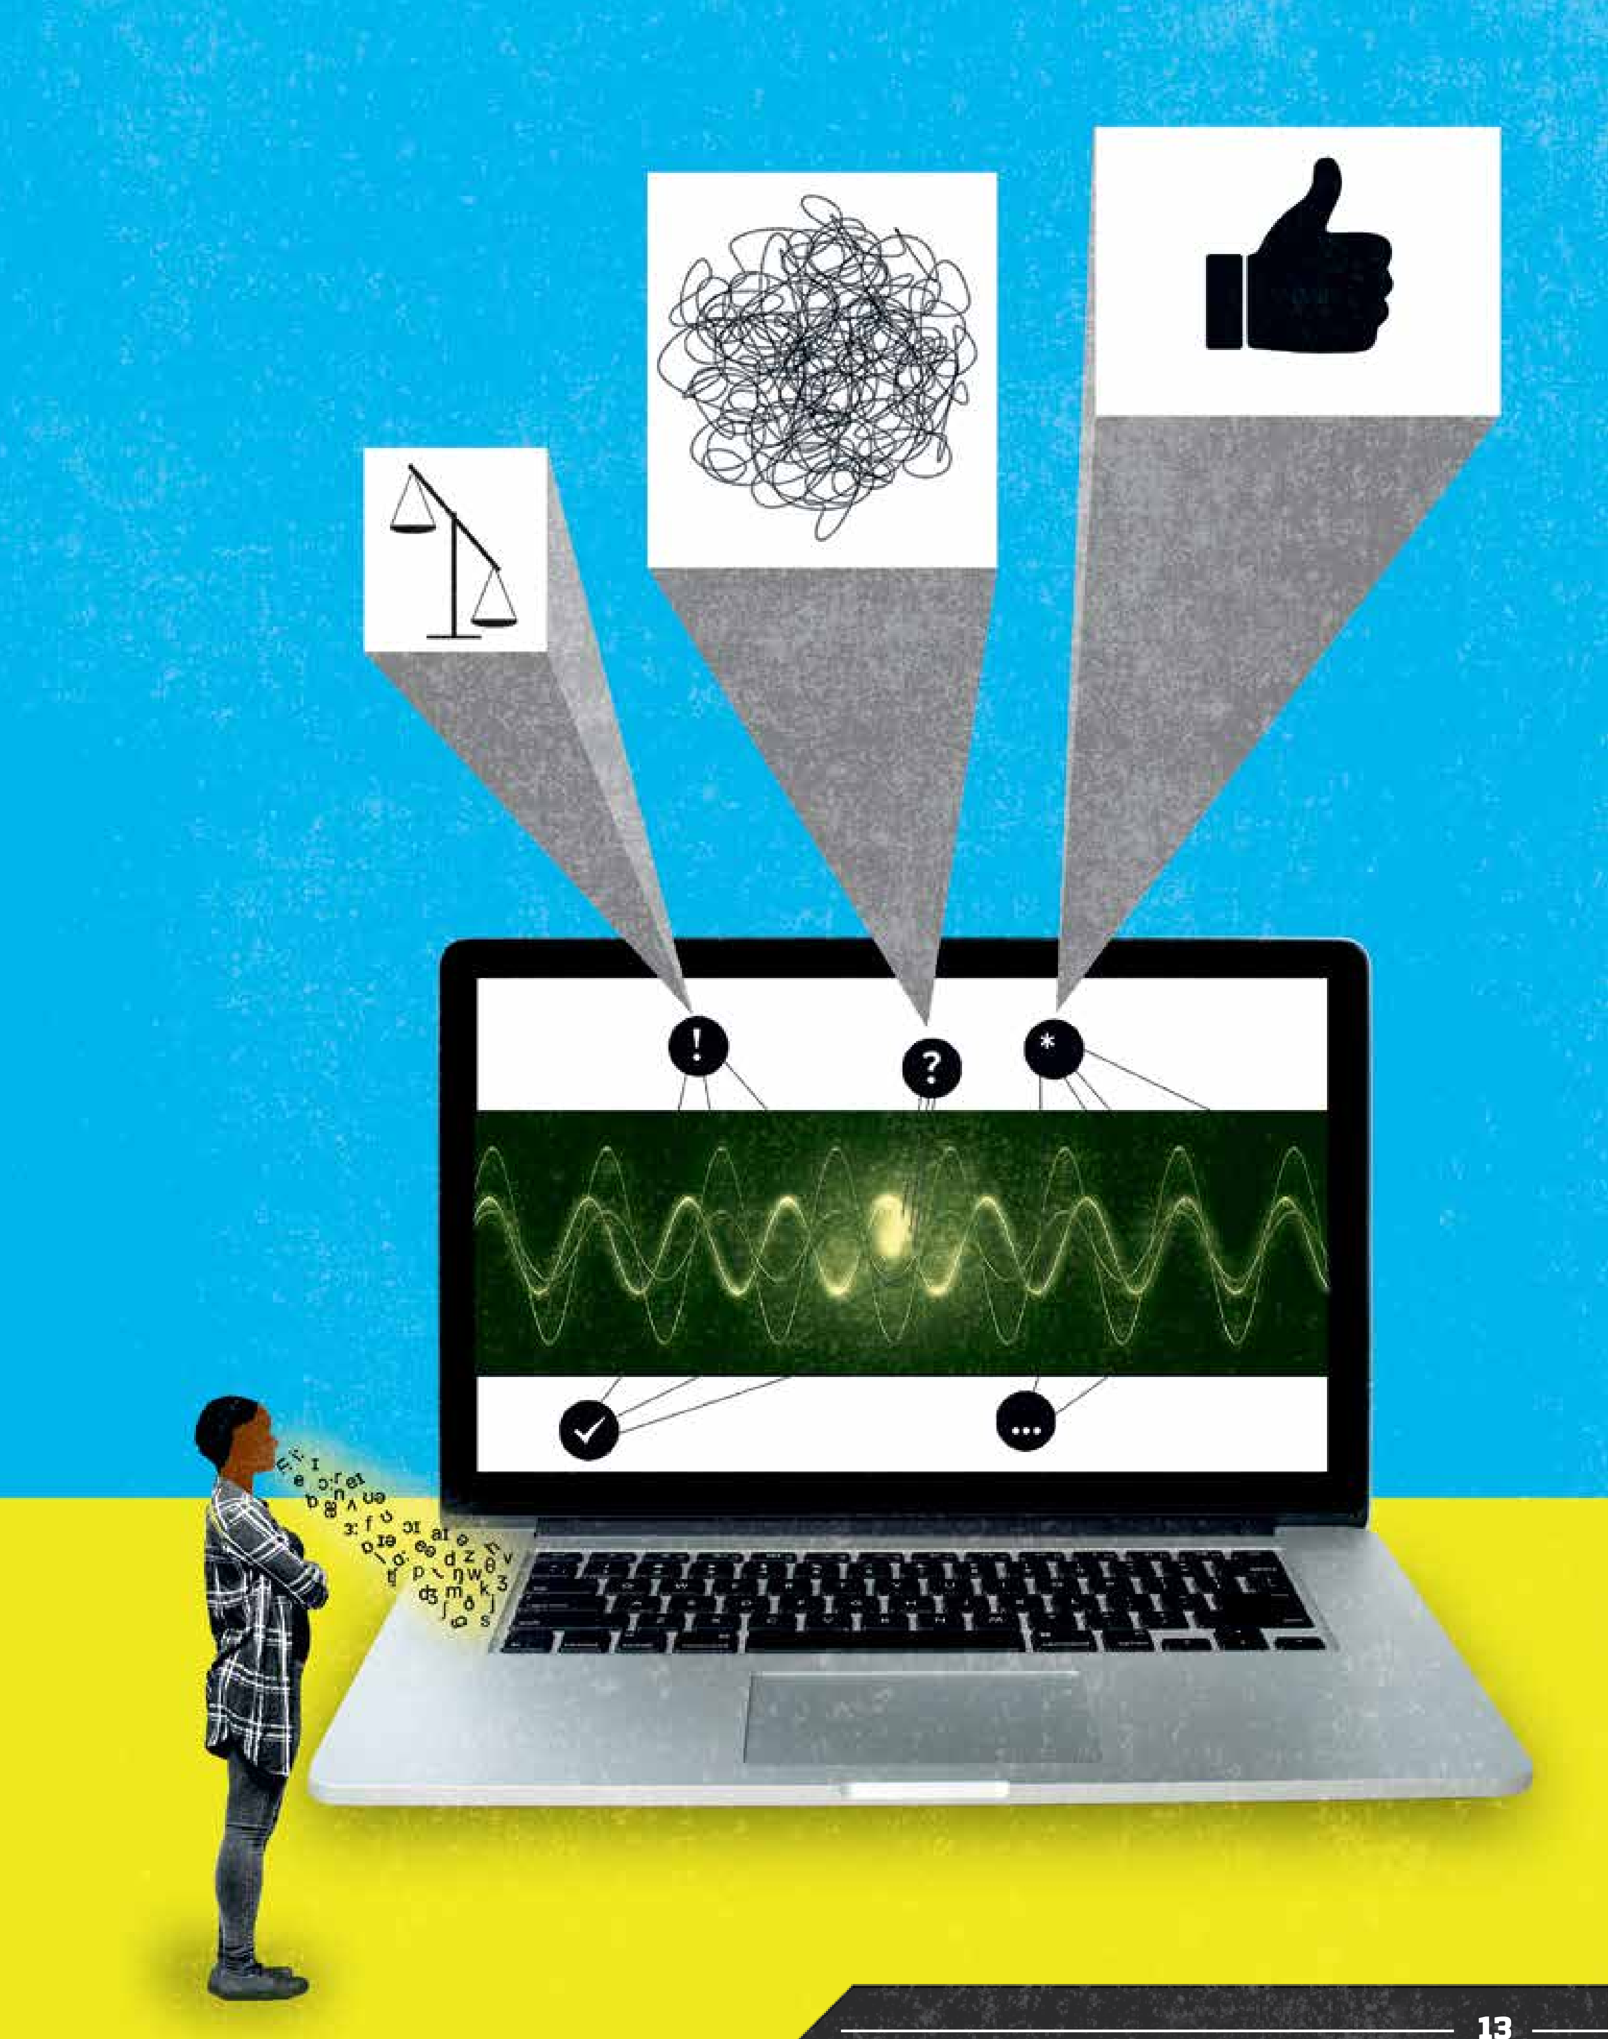

What's in a name? Sometimes, more than might be expected. Assistant Professor of Computer Science **Anjalie Field** has shown that something as seemingly innocuous as people's names can offer insight into how artificial intelligence and machine learning can get it wrong, and by extension, do wrong.

Last year, Field—a member of the Center for Language and Speech Processing—was part of a team of researchers investigating the possible application of natural language processing (NLP) to assist the nation's network of child protective services agencies in better serving and protecting vulnerable children at risk of abuse or neglect. NLP is a form of artificial intelligence (AI) that processes large datasets of normal human language using rule-based or probabilistic machine learning to contextualize and understand written communication.

Protective services agencies' case workers typically take careful and extensive notes related to every family that enters the system. Typically, those notes are not analyzed in any systematic way to learn how better to respond to and possibly prevent child abuse. In the social services community, there is tremendous interest in using NLP and AI to help better manage the more than 3.6 million calls that come into child protection agencies across the United States each year.

Field's research team applied NLP tools to 3.1 million contact notes collected over a 10-year period in an anonymous child protective services agency. They found that the NLP model did a poorer job at recognizing the names of African American individuals, which could potentially include identifying relatives who might support or shelter endangered children, than it did with white individuals. "Here is an unlooked-for area where machine learning can amplify what is already a well-documented racial bias in child protective services without anyone even recognizing the problem," Field says.

As AI and machine learning technologies increasingly come to inform decision-making tasks—and in some situations make decisions—there is increasing concern over issues of justice, equity, and fairness. At the Whiting School, faculty researchers are engaging with the fundamental issue of fairness in artificial intelligence across a wide spectrum of potential uses and misuses. These efforts range from unleashing the newfound power in neural network–based computing to search for and identify existing biases, to combatting the known tendency of machine learning systems to learn and propagate unfair practices.

Investigate truth and fairness in machine learning and you quickly discover there is more than one aspect to the challenge. Field, whose research focuses on the ethics and social science aspects of NLP, says her work follows three broad directions. "One of them is computational social science: How can we use language analysis and automated text processing to identify social injustices that already occur in society?" This approach employs machine learning technology as a tool for identifying bias. For instance, Field has used NLP to investigate the language Wikipedia articles use to portray different groups of people. "Do articles about women emphasize their personal lives more than their careers, whereas articles about men talk more about their careers? When you employ these tools to look across all of Wikipedia and compute statistics around the data, you can start to see these kinds of patterns and biases," she says.

A second line of her research focuses on building technology to address known issues of existing bias, particularly in public service domains. As part of her work looking at child welfare cases, she used NLP to review thousands of case records, categorized by race, for the frequency of certain words and terms, to identify persistent institutional bias. She found differing language, but not enough to provide conclusive evidence of racial bias. "This is sort of natural language processing for the social good," is how she describes it.

The third direction looks at potential ethical issues within AI itself. "Some of that is directly looking at model bias—is it going to favor certain groups of people more than others?—but also issues like access to technology, privacy, and data ownership." In this broader subject area of bias, it is the data itself that is typically of greatest concern, but Field emphasizes that bad data should never permit bad results. "I think there's often a narrative in AI that 'the data is biased so it's not our fault/there's nothing we can do about it,' but that's not true," she says.

Mark Dredze, John C. Malone Professor of Computer Science, a pioneer in using AI tools to gain insights into public health challenges ranging from suicide prevention to vaccine refusal and from tobacco use to gun violence, has also found how bad data can work to amplify bad outcomes. He says early missteps in machine learning algorithms—such as when Microsoft's initial release of a chatbot called Tay had to be shut down after just 16 hours for spewing obscene and racist language it picked up online—highlight the dangers of working with bad data.

"The problem is that if we just accept the data and say to the machine, 'Learn how to make decisions from the data,' when there are biases in the data, or biases in the process of how we learn from data, we will produce biased results. That is the essence of the problem of fairness in machine learning algorithms," Dredze says.

"Do articles about women emphasize their personal lives more than their careers, whereas articles about men talk more about their careers? When you employ these tools to look across all of Wikipedia and compute statistics around the data, you can start to see these kinds of patterns and biases."

Anjalie Field

This applies to his own research, some of which involves scanning the web to understand how people turn to the internet for medical information and what innate biases they are likely to encounter there. In a recent paper he co-authored, researchers set about evaluating biases in context-dependent health questions, focusing on sexual and reproductive health care queries. They looked at questions that required specific additional information to be properly answered when that information was not provided. For instance, “Which is the best birth control method for me?” has no single correct answer, as it depends on sex, age, and other factors. Dredze and colleagues found that large language models often will simply provide answers reflecting the majority demographic, as, for example, suggesting oral contraceptives, a solution only available to women, while neglecting to include the use of condoms. This kind of built-in bias is a special concern for individuals who turn to the web as a replacement for traditional health care advice since misinformed answers have potentially detrimental effects on users’ health, Dredze says.

This and his earlier pioneering work in “social monitoring”—employing machine learning to gain understanding from text published on social media sites—has led him to focus not just on the raw data, but also on how people use the web.

“I would describe this as a more holistic approach, where we’re actually building systems and paying attention to how people interact with those systems,” he says. “Where did our data come from? How did we collect it? I have to care about these issues when giving data to the algorithm and figuring out what the algorithm does. But then I also need to account for the fact that a human will interact with us. And humans are going to have their own biases and issues. So maybe it’s not just that the system is biased or unbiased, but it interacts with someone to create a different kind of bias.”

The Data Decides

So how do you make sure AI and programs based on NLP such as ChatGPT are fair and free of bias? To understand the nature of the challenge presented requires grasping a fundamental paradigm shift that has occurred. “What’s different is that in the traditional ways computers have been used, the programs themselves were the issue—the code made the decision,” explains cybersecurity expert **Yinzhi Cao**, assistant professor of computer science and technical director of the JHU Information Security Institute. “For AI, the most important thing is the data, because AI learns from the data and makes decisions from what it learns.”

It is that capability of making autonomous decisions that has the potential to be especially unfair, Cao says. His research has recently focused on security, privacy, and fairness analysis of machine learning systems. He notes that these three concerns can overlap in surprising ways.

Cao was one of a group of researchers to successfully overcome safeguards on two of the better-known text-to-image generative models that use AI to create original images from written prompts. Typically, these art generators are designed with filters to block violent, pornographic, or other objectionable content. But Cao and colleagues showed that the right algorithm could be used to bypass filters and create images that are not simply unsuitable but also could be used to defame or malign individuals “like a politician or famous person being made to look like they’re doing something wrong,” Cao said in a press release announcing the results of the team’s research efforts.

Cao also conducts research that instructs AI in medical image analysis to augment medical training and diagnosis in the detection of Lyme disease, a project he has worked on in cooperation with colleagues at the Applied Physics Lab and the School of Medicine. An early symptom in 70% to 80% of people with Lyme disease is the appearance of a distinctive “bull’s-eye rash,” which is usually a single circle of inflamed skin that slowly spreads from the site of the tick bite. In fair-skinned people, and especially in younger people with unblemished skin, it is easily detected. For clinicians, however, the most useful examples are found in skin types where the contrast is not as distinctive and readily apparent. “The final goal may be that we will be able to ask AI to actually perform the diagnosis, but at this stage, we are using it to generate training images that overcome disparities in age and race and gender,” he says.

“The problem is that if we just accept the data and say to the machine, ‘Learn how to make decisions from the data,’ when there are biases in the data, or biases in the process of how we learn from data, we will produce biased results. That is the essence of the problem of fairness in machine learning algorithms.”

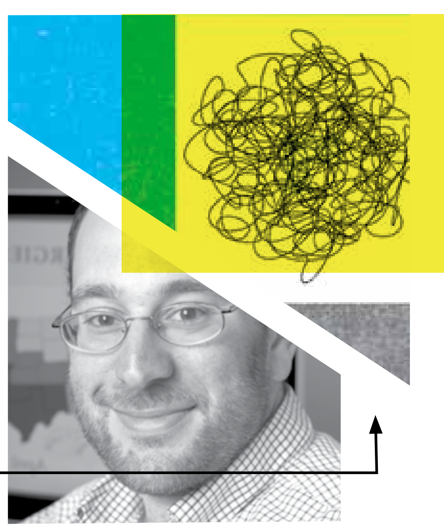

Mark Dredze

Algorithms Are Everywhere

Making AI decisions understandable may first require overcoming a larger challenge. The term algorithm itself raises math anxiety among many since the word is poorly understood. But **Ilya Shpitser**, the John C. Malone Associate Professor of Computer Science whose work focuses on algorithmic fairness in datasets of all types, points out that algorithms are the basis for everyday decision making among many—even if they don't know it.

"When doctors diagnose, when judges set bail, they have a sequence of steps they've learned that's considered reasonable," he says. "Regardless of how they think of it themselves, they are using algorithms, because judicial decisions and diagnosis cannot be arbitrary; they better be systematic. The fact that similar cases are decided in similar ways: That's what an algorithm really does."

Most important for a judge to appropriately set bail or a doctor to accurately make a diagnosis is the need for good, fair, accurate information. In algorithmic decision-making, it all comes down to the data. And in an imperfect world, for any decision, there will be good data, there will be bad data, and most vexingly of all, there will be data we simply don't have.

"Any person who works in actual data, real data of any kind, has missing data in their sets, that's just how it is, basically," says Shpitser, who cites a common example of data collection in electronic health records, where missing data can be due to a lack of collection such as a patient was never asked about asthma. Or it could come from a lack of documentation, as when a patient was asked about asthma but the response was never recorded in the medical record. "Lack of documentation is particularly common when it comes to patients not having symptoms or presenting comorbidities." In these cases, rather than recording a negative value for each potential symptom or comorbidity, the missing data fields are left blank and only the positive values are recorded, which skews the value of the data, says Shpitser. "This makes it essentially impossible to differentiate between the lack of a comorbidity, the lack of documentation of a comorbidity, or the lack of data collection regarding the comorbidity."

One of the central challenges in creating fair and accurate algorithms then becomes devising sound methods of correcting for data recorded incompletely, incorrectly, or not at all. "I work on data being screwed up," is how Shpitser describes it. Along the way, he has demonstrated in his research that it is at least theoretically possible to "break the cycle of injustice" (in which variables such as gender, race, disability, or other attributes introduce bias) by making optimal but fair decisions. His research employs the methodology of causal inference, which he describes as "methods to adjust for incomplete, bad, or missing data to allow reliable and fair inferences to be made."

Fairness Isn't an Oracle Concept

In the past few years as AI systems have caught increasing media attention, highly public machine learning misfires have brought scientists and engineers a deeper awareness of both the importance and the difficulty of designing and implementing systems equitably.

"For a long time, we said, 'Look, the algorithms are math, and math is math,'" says Dredze. "It was, 'Let's throw the math at this, and it comes up with what it comes up with.'" That attitude no longer applies. "I think we've learned a couple things. One is that the math might be math, but the data is not the data: It always has some kind of bias in it. And we need to do something about that."

But that may only be the beginning. "The other thing we're learning—and maybe there's a little controversy to this—is that math isn't just math. Math always has some assumptions to it. The models that you pick always have some assumptions, and for a variety of reasons we might favor certain models that do better on some groups. And so it's not only a matter of the data; it's also a matter of the models we build. How do we make our models aware that fairness is a thing? How do we build into the models some measure of fairness?"

Which points ultimately to issues that transcend both the data and the math.

"Fairness isn't an Oracle concept," notes Dredze. "If you've got kids and you've ever tried to give them

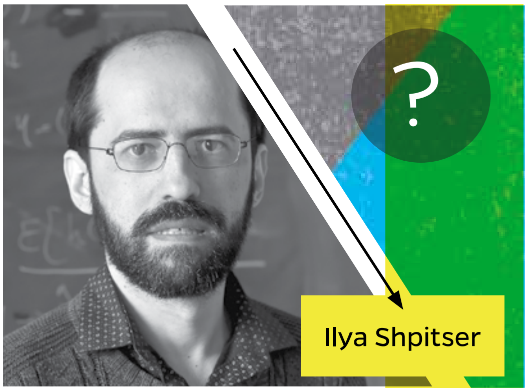

Ilya Shpitser

"My take on algorithmic fairness research is that it's not our job as researchers to decide what fair is. I think the discussions of fairness need to be discussions in the public square."

"The final goal may be that we will be able to ask AI to actually perform the diagnosis, but at this stage, we are using it to generate training images that overcome disparities in age and race and gender."

Yinzhi Cao

anything, they complain about fairness, right? And you end up telling them that life isn't fair. Fairness is subjective: that person got a smaller piece of cake, but they had a frosting flower, and you didn't get a flower."

It is, he points out, tremendously challenging to formalize concepts of fairness. "We can build that into our models and train them to be aware. But who decides what the right definition of fairness is? Think about the most controversial issues in society, like college admissions. Both sides of affirmative action and college admissions are insisting we need to be fair, but they have opposing views as to what that means," he said.

All of which suggests that creating algorithms for a fairer world in the end will not be the purview of computer scientists alone.

"My take on algorithmic fairness research is that it's not our job as researchers to decide what fair is. I think the discussions of fairness need to be discussions in the public square," says Shpitser. "As an American citizen, I have my own opinions of what policies we should follow, but that's a different path than wearing my hat as a researcher on algorithmic fairness. In other words, computer scientists are best suited to be the implementers, not the deciders, in notions of what is fair."

He continues: "Whenever I give talks about this, I always get questions that try to push me into being some kind of priesthood that decides for people what fairness criteria to use. I really don't think that is our job. I have as much grounds to advocate for a particular definition of fairness as any other citizen, but the fact that I work in algorithmic fairness doesn't give me a special advantage. I'm using modern tools, but the questions themselves are much older than that."

Overcoming 'Ethical Debt'

CATHY PETROZZINO '80
Mathematical Sciences

How fast is too fast? **Cathy Petrozzino** '80 believes that in the realm of AI and large language model chatbots, this is a question we need to be asking. As a principal in cybersecurity, privacy, and AI ethics at the MITRE Corporation, a public interest nonprofit in Bedford, Massachusetts, that operates federally funded R&D centers on behalf of the government, she is paying close attention to the speed and ubiquity with which these tools are being adapted and employed. In particular, Petrozzino doesn't believe an organization's good intentions alone are sufficient to identify and manage ethical risk in AI.

Petrozzino sees the integrity and reliability of data used in AI as a defining limitation, describing the challenge as a form of "ethical debt." She borrows from the concept of "technical debt," which compares releasing first-time computer code that has not been fully tested and corrected to going into debt. Sending out code with technical debt speeds development and can be a good thing—so long as it is "paid back" promptly to ensure the integrity of the code base.

"Ethical debt is the same kind of concept," she explains. "It's when you're more interested in putting out a solution quickly, which is often the business driver—to be first to market. But you haven't yet looked at your solution from an ethical perspective and you end up having a situation where there are biases or other problems within the data or AI solution that haven't been addressed. And so, there's ethical debt."

But there's a crucial difference. Technical debt is owed by developers to customers who bought the product. Ethical debt falls not on customers, but on individuals who have no concept that they're being treated unfairly because an algorithm invisible to them was used and the humans accountable for developing, deploying, and sustaining this algorithm did not adequately address ethical risk. "Often, the unfortunate truth is that it's the people—particularly individuals or groups who have been historically marginalized—who pay the price as a result."

Petrozzino describes herself as optimistic about finding ways to overcome ethical debt. "I look at this as a classic risk management problem. There are much more risk identification and management activities, and more thought about how we manage the risk today than in the past, thanks partially to ChatGPT. So I am encouraged by that."

–Mike Field

Trauma Response

Many deaths from traumatic injuries that cause uncontrolled internal bleeding are potentially survivable if an unseen injury is diagnosed and treated promptly. **Lidia Al-Zogbi**, a fifth-year doctoral student of mechanical engineering in Axel Krieger's, PhD '08, IMERSE (Intelligent Medical Robotic Systems and Equipment) Lab, has designed a trauma robot that aims to lower the number of deaths from undiagnosed internal bleeding. The autonomous system can alert doctors to a hemorrhage via ultrasound and initiate diagnostics to stabilize patients en route to the emergency department.

"This will help ensure that patients make it to the hospital, where doctors who have been alerted to their condition are waiting," she says.

—JONATHAN DEUTSCHMAN

Lidia Al-Zogbi

PITCH-PERFECT VOCALS

Once a solution for amateur, off-key singers on karaoke nights, Auto-Tune and other voice tuning software have progressed from the club to the recording industry but never beyond sounding artificial.

Now, Whiting School researchers have made significant improvements to traditional voice tuning software's previous capabilities. This innovation—called Diff-Pitcher—seamlessly corrects out-of-tune singing while maintaining the original vocal timbre and naturalness, expanding its possible applications beyond entertainment and the music industry and into the health care setting.

"Diff-Pitcher is a generative deep neural network that takes pitch-correction technology to a new level. Its precision and control not only help musical artists and producers but also open new possibilities in voice rehabilitation and assistive technologies," says team member **Jiarui Hai**, a PhD student in the Department of Electrical and Computer Engineering, who worked on the project with lead researcher **Mounya Elhilali**, a professor in electrical and computer engineering.

"[The results sound] really natural, and unlike in older ways of fixing pitch, we can still regulate how high or low the voice goes," Hai says. "The technology could revolutionize treatment for a spectrum of speech-related disorders, offering valuable support for post-laryngectomy patients and contributing to the voice rehabilitation of stroke victims."

— DINO LENCIONI

‘CREATE AN AI NEURAL NETWORK CHIP’

Graduate student **Michael Tomlinson** and undergraduate **Joe Li** have tapped into the promise of natural language prompts and ChatGPT4 in a pioneering new approach to creating a neuromorphic spiking neural network chip for machine learning and AI.

Through step-by-step natural language prompts to ChatGPT4, starting with instructions to create the spiking behavior of a single biological neuron and then linking more to form a network, they generated a full chip design that could be fabricated to build a larger network chip that abstracts the functionality of neurons in the human brain.

Tomlinson believes this is the first AI chip that is designed by a machine using natural language processing.

“The process can be likened to giving instructions to a computer, such as ‘Create an AI neural network chip,’ and the computer generates a file containing the specifications and design for manufacturing that chip,” says Tomlinson, who worked closely on the project with Li and with **Andreas Andreou**, a professor of electrical and computer engineering who is also co-founder of the Center for Language and Speech Processing and is affiliated with the Kavli Neuroscience Discovery Institute and JHU’s new Data Science and AI Institute. The team’s design appears on the preprint site arXiv.

Tomlinson explains that their design was simple as a proof of concept. The chip’s final network architecture is a small silicon “brain” with two layers of interconnected neurons. It has a flexible programming system, featuring an 8-bit addressable weight system to program the chip and fine-tune the chip’s behavior. This customization is facilitated through a user-friendly interface, akin to a remote control.

Before manufacturing, the team validated the chip design extensively through software simulations to ensure functionality and address any issues. Then they submitted the design electronically to the Skywater “foundry,” a chip fabrication service in Minnesota, where it is currently being “printed” using a relatively low-cost process.

“While a small step, this shows AI can create advanced hardware to accelerate the development of technology,” says Tomlinson, noting that neural network chips could soon power energy-efficient, real-time machine intelligence for next-generation embodied systems like autonomous vehicles and robots.

Tomlinson’s research interests are in building systems to solve the computational problems posed by today’s cutting-edge algorithms, particularly for AI. “These algorithms are getting more expensive to compute,” he says. “I would like to rethink how we architect our computing systems to more efficiently implement these algorithms and solve problems that people care about.”

—DL

> “Eventually, AI will be used in nearly all industries. Within 15–20 years, it will mimic humans in numerous complex tasks, often with better accuracy.”

— **Alok Aggarwal**, PhD ’84, is the author of *The Fourth Industrial Revolution and 100 Years of AI (1950–2050)*, and CEO and chief data scientist of ScryAI, a San Jose–based analytics firm that develops AI software solutions for the financial, technology, and health care industries. ScryAI is named after an early 16th-century verb that means “foretelling the future using a crystal ball.”

In 2014, Aggarwal launched the Silicon Valley startup, which now has 320 employees and offices in Delhi-Gurgaon, India. “We’re utilizing algorithm-based methods to provide intelligent document processing,” he says. The company’s signature product, Collatio, inspired by the word “collation,” extracts financial, legal, and operational Excel spreadsheets and PDFs for end users. Typical clients for the quick and accurate organizational software include financial institutions providing business loans.

ScryAI offers four additional software programs to detect anomalies, predict operational risks, identify fraud, and extract data. Its Data Flow Mapping program utilizes more than 60 proprietary algorithms to reverse engineer previously programmed processes.

Before founding his current enterprise, Aggarwal was for 13 years co-founder and chairman of Evalueserve, a research and analytics company where he pioneered the concept of knowledge process outsourcing. He previously held various key leadership roles at IBM for 16 years, including founding and directing the IBM India Research Lab, which was affiliated with the Indian Institute of Technology Delhi, where he earned his BTech degree in 1980.

— NIKOLAS CHARLES

In AI We Trust

ILLUSTRATIONS BY Patric Sandri

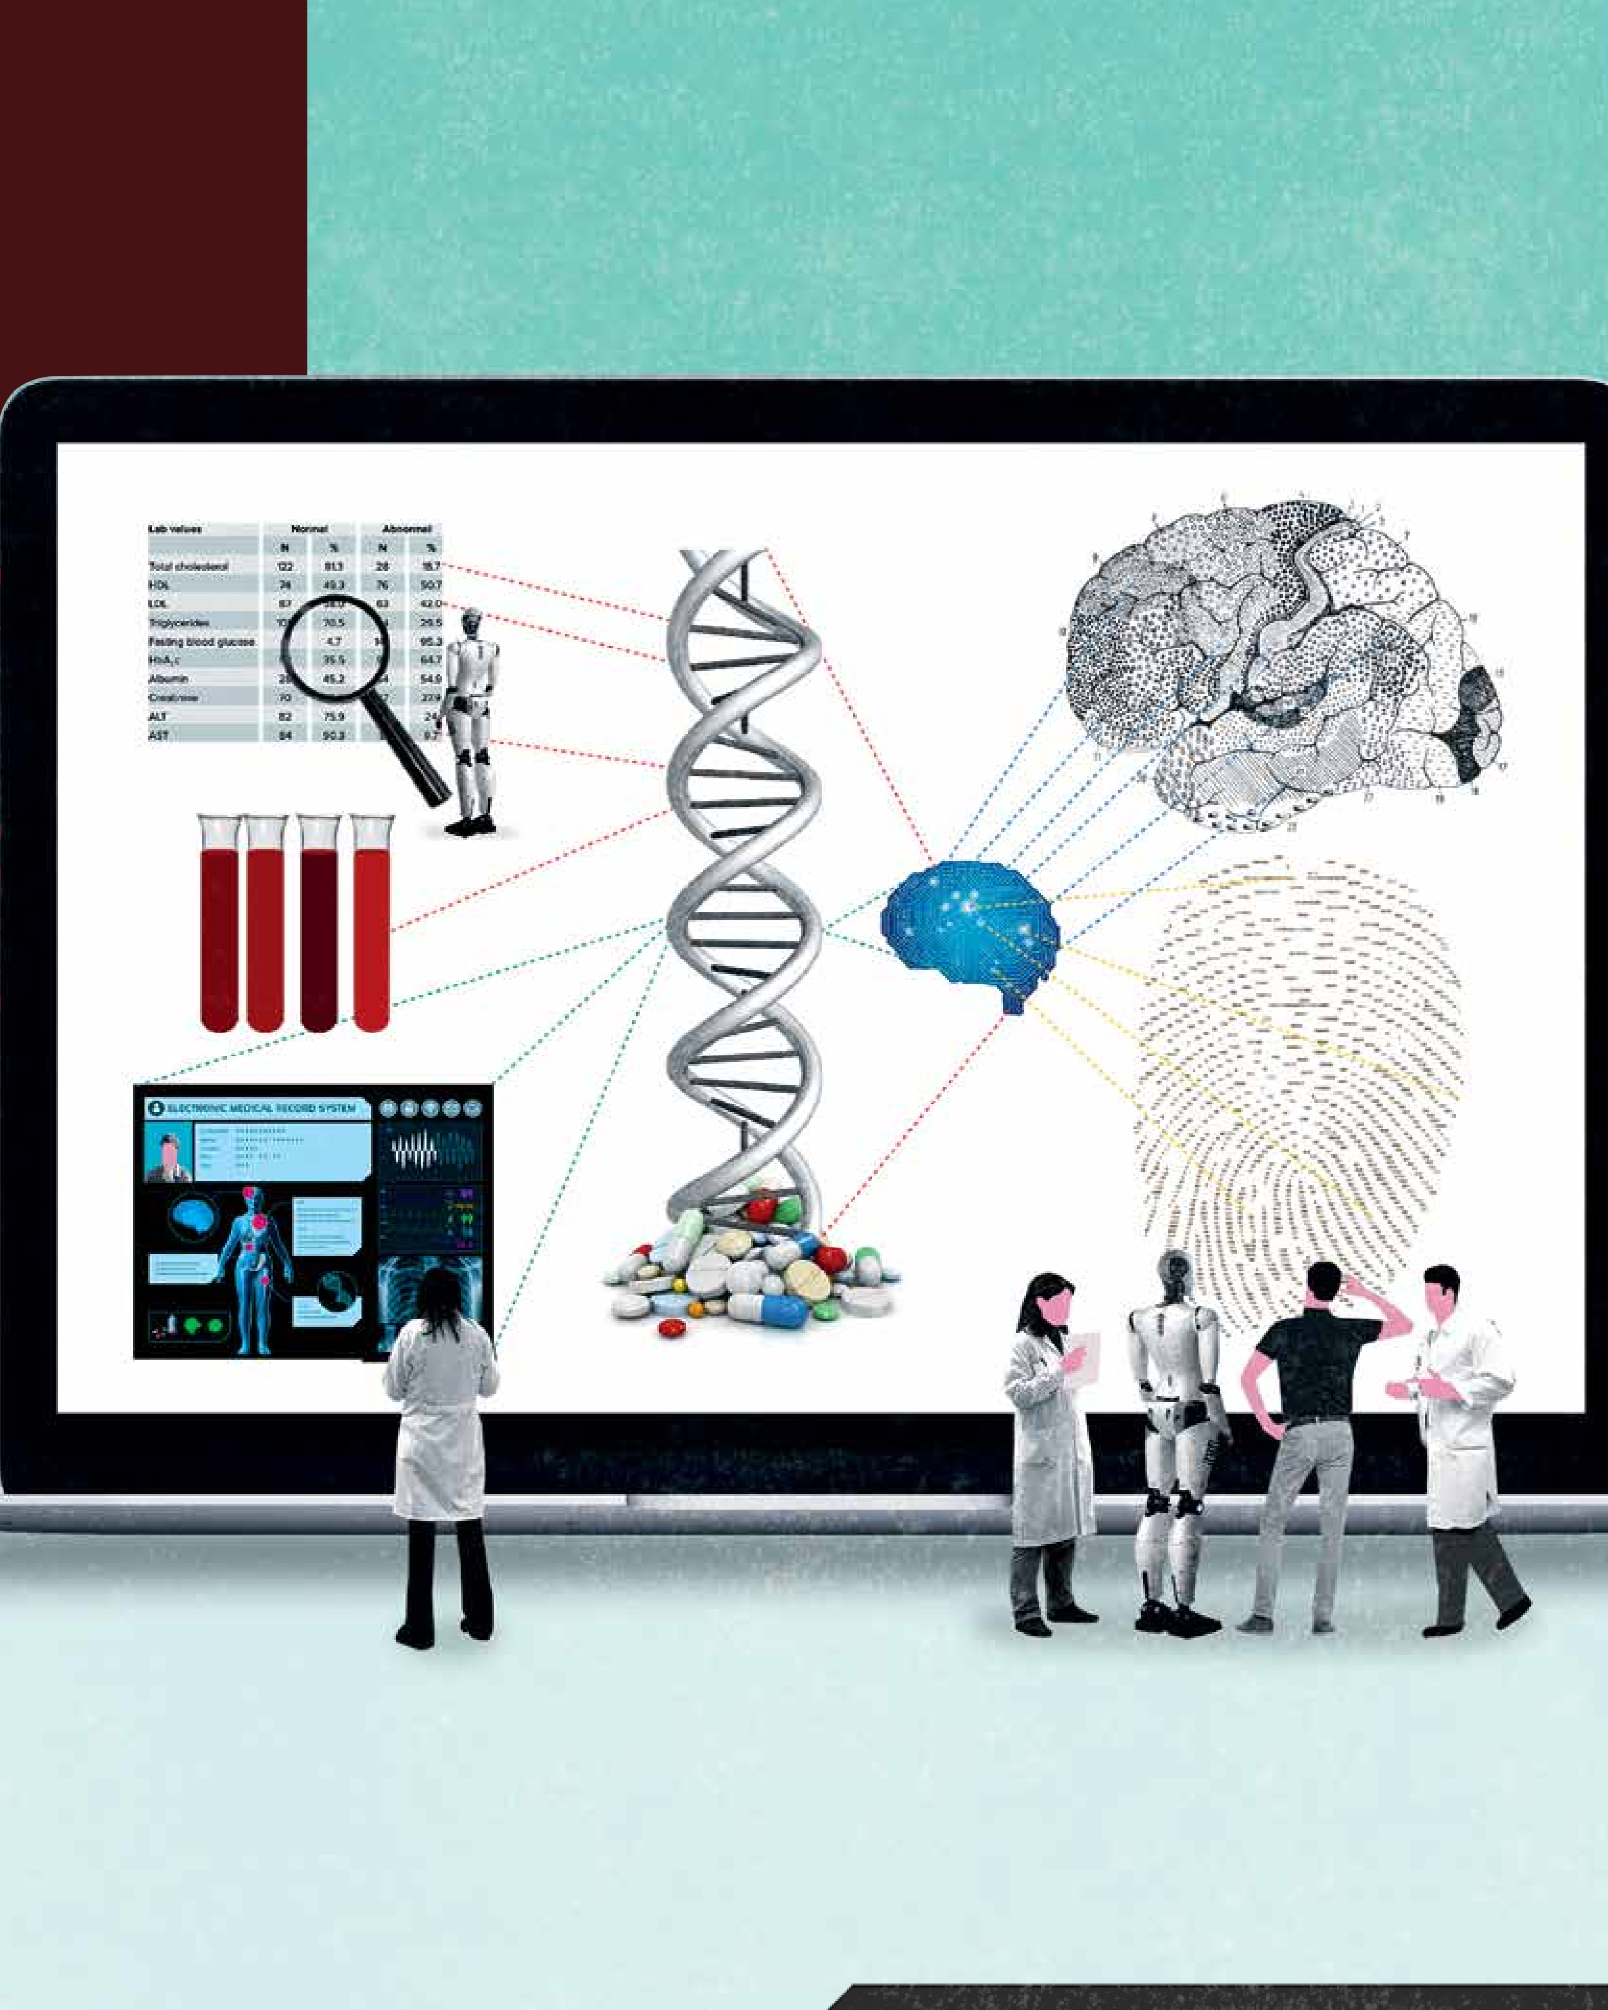
Lab values
Normal
Abnormal
Total cholesterol
HDL
LDL
Triglycerides
Fasting blood glucose
Albumin
Creatinine
ALT
AST
ELECTRONIC MEDICAL RECORD SYSTEM

We increasingly rely on AI models in our daily lives—from traffic navigation and shopping apps to AI-informed care decisions made by our doctors.

Given their ubiquity and influence, how and why should we trust these decisions? Can we be certain the models' predictions are free of biases or errors?

To explore these and other timely issues that academic researchers are grappling with—Can and should AI be regulated? How have rapid AI advances changed teaching?—we brought together a panel of AI experts from across the Whiting School and tapped AI researcher **RAMA CHELLAPPA** to lead them in a wide-ranging exchange. Highlights from their discussion follow.

DANIEL KHASHABI
Assistant professor, Computer Science Center for Language and Speech Processing

RAMA CHELLAPPA
Bloomberg Distinguished Professor in electrical and computer engineering and biomedical engineering; Co-interim director, Data Science and AI Institute

MATHIAS UNBERATH
Assistant professor, Computer Science Laboratory for Computational and Sensing Robotics, and Malone Center for Engineering in Healthcare

NATALIA TRAYANOVA
Murray B. Sachs Professor, Biomedical Engineering; Director, Alliance for Cardiovascular Diagnostic and Treatment Innovation

The 'Black-Box Problem'

RAMA CHELLAPPA: When we go to a doctor's office and we see diplomas from Harvard or Hopkins, we think the doctor knows what they are doing, and we trust them. But if the doctor is [employing] AI, we want to know where AI learned to diagnose illnesses and recommend therapies. When we have a physician, a patient, and AI working together, that requires complete trust. So, interpretability is a very key component to trusting the AI-based decision.

When it comes to artificial intelligence, how does the "black-box problem" impact your work?

NATALIA TRAYANOVA: In my center—the Alliance for Cardiovascular Diagnostic and Treatment Innovation—our work is very clinically translational. Pretty much everything that we develop, our goal is to immediately bring it to the patient's bedside. We work very closely with clinicians. We are fully embedded in the Department of Cardiology, particularly interventional cardiology.

In medical applications of AI, it is really important that we can learn how decisions are made. Black-box algorithms can restrict the clinician-patient relationship. The medical advice the doctor provides depends on medical reasoning and clarification, which is not available from black-box algorithms. When a black-box AI algorithm provides only a decision and not a justification for that decision, patients are deprived of an understanding of the underlying medical problems and why the decision is made. This can hurt the clinician-patient relationship and erode trust in the medical system. It's on our shoulders to develop AI applications that create and maintain that trust.

MATHIAS UNBERATH: This black-box problem requires nuanced discussion. The AI community is proposing explanations as a possible answer, but I think explainable AI may not always be the preferable solution—and it may not even be necessary. Rigorous testing and evaluation to help elevate trust in AI at the technology level together with broader education will also contribute to an answer.

But of course, achieving trustworthiness now is important. To contribute to this goal, we are trying to understand a few things. One is to better understand the datasets being used to develop these algorithms in hopes that understanding data-inherent biases will help us create models that don't perpetuate [those biases]. Insights from causal reasoning as well as synthetic data can play an important role in this endeavor.

The other is about how we deploy these models. If I train a model, will it be self-sufficient, and involve no human interaction? In some cases, this might be happening, but in other high-stakes environments, such as medicine, we might not want to run the model without human interaction. Then, we don't need to think exclusively about the model and its performance, but about the model and its possible explanations in the context of human interaction. That is a much more complicated problem.

DANIEL KHASHABI: I want to second what Mathias said about making sure that users are part of this analysis. I don't think there will ever be an AI system that has responses interpretable to every human, so we need to define interpretability with respect to certain audiences.

For example, people trust airplanes, not because we can "interpret" them, but because we trust that someone who's an expert in airplanes can interpret their functionality.

The same logic will likely hold also for various applications of AI. We should aim to build interpretable AI applications with high accuracy in, for example, the medical domain for experts in that domain. Increasingly, people will trust that others can interpret those results.

MATHIAS UNBERATH: What Daniel said about tailoring explanations to people is interesting. Right now, there seem to be very different philosophies in how we do "explainable AI."

There is the philosophy driven by people in the computer science/computer vision/machine learning AI community who are very interested in building novel explanation techniques that satisfy the needs of computational demand such as factual accuracy or fidelity. These people focus on computation first, with users and humans often being omitted or limited to an afterthought.

And there's the human-computer interaction community that publishes in completely different venues. Their papers look and feel completely different, and they don't over-prioritize technical feasibility, especially early on, in favor of a more human-centered approach. They are focusing on tailoring a specific message to a specific user and measuring that. These two communities don't really connect.

NATALIA TRAYANOVA: To connect to Daniel's example of why people trust airplanes, patients don't question why they are given aspirin. They don't understand the mechanism by which aspirin works, but they take it because it's so commonly prescribed and has undergone so many clinical trials. I think this is something the AI community can look forward to.

Danger Ahead?

Let's move on to talking about the risks of AI.

DANIEL KHASHABI: I will briefly touch upon risks that keep me worried about the future of AI.

AI thrives on data, and with the push to make AI systems more accurate, there's a lot of appetite for more data, and especially user data. The private sector is going to continue to collect increasingly more personalized data. Legislation in this country is so passive. That basically means that citizens here need to be vigilant, monitoring emerging applications and watching for anything that could possibly go wrong as a result of overuse or overreach of personalized data.

Relatedly, I'm quite worried about the personalization of propaganda. In 2016, when we had an election, there was plenty of evidence that social media platforms were used to create propaganda to interfere with the election. Now that AI is being industrialized, it's possible to personalize this propaganda to an individual's mindset and you can frame your arguments to an individual's world model, values, and their most personal weaknesses, to sway their opinion. This is a really dangerous weapon for attacking a democracy.

MATHIAS UNBERATH: I have similar concerns, including that it's not always transparent what the objectives are that specific algorithms are optimized for.

Think about bail decisions or credit approval. The values or the reasons that go into making such decisions are very complicated and quantifying them in a way that can be optimized computationally via empirical risk minimization is not trivial. And there's an additional risk that we can introduce

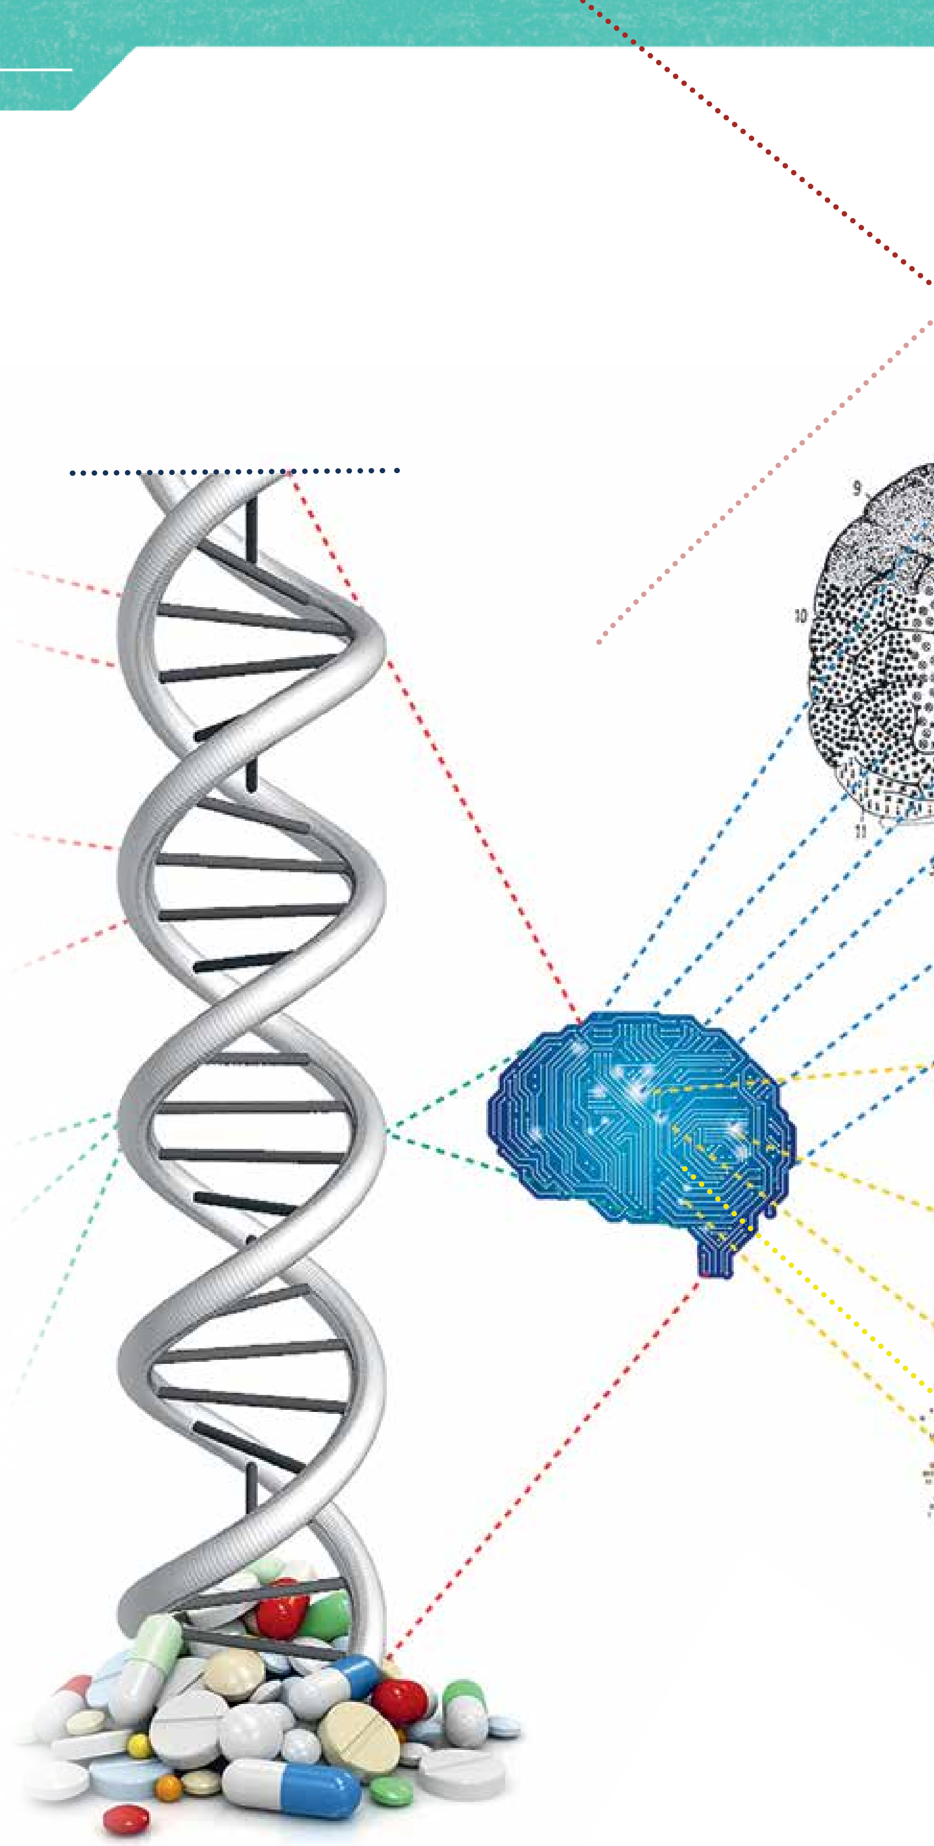

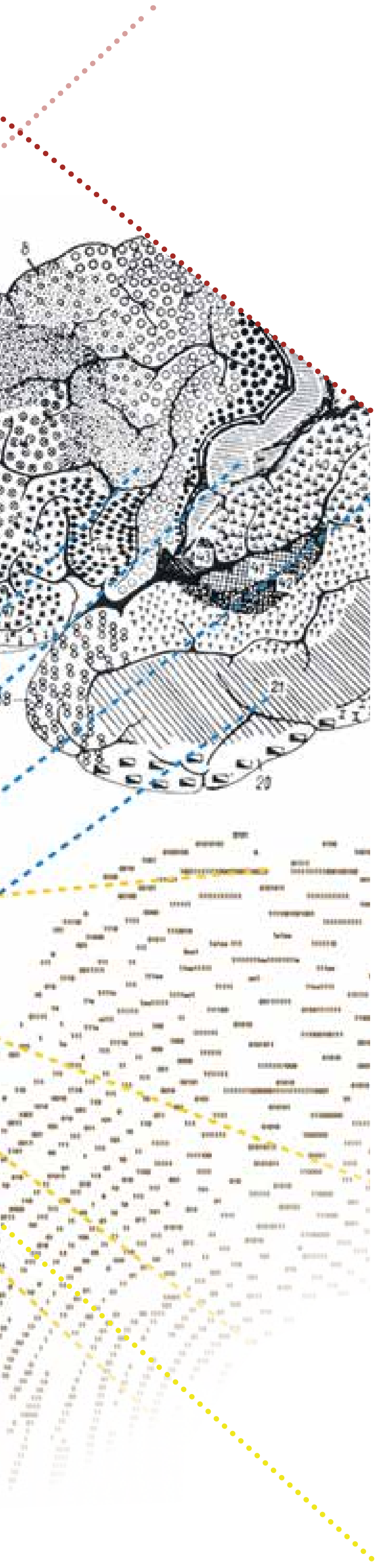

malicious behavior or exploitative intent by simply optimizing models to prioritize specific scenarios. Combine this with scalability of automated approaches, and poor decisions might be made at unprecedented scales.

Some say that we don't need to fear AI because we get to assign power to it, but I wholeheartedly disagree with this and that is because there is "creep." The presence of AI affects not just whether I consciously do or do not use it. AI can affect people's experiences subtly. For example, when it's used in areas like content moderation, it can change the landscape. On the individual level, the effects might be marginal, but at population level, they might tip the scales in favor of one policy or outcome, and that's power that nobody consciously assigned to AI.

RAMA CHELLAPPA: As far as risks, I always talk about domain shift—you train AI using one kind of data and you change the data, and the performance becomes unpredictable. This happens in pathology, when different labs have different kinds of images, with slight changes.

I also worry about attacks on AI systems and bias—whether a system is fair to every group.

AI's ability to generalize so it works for everybody makes it robust to attacks, but because people will try to attack any software system, we have to continuously monitor it. I think of it as a precocious child who will let the water run over in the kitchen sink to understand hydrology. We have to tell the kid to learn things properly and be useful.

MATHIAS UNBERATH: I agree with you. I am incredibly excited, and AI is a large part of my research agenda. I just think there's a lot of hype right now [and what's getting lost] is a more cautionary approach—ensuring AI is scalable and that we can sustain its ethical deployment. I think those voices are a bit suppressed.

Should AI be regulated ... and is regulation even feasible?

MATHIAS UNBERATH: I am not a regulatory expert, so I don't think I sufficiently understand the regulatory powers and opportunities to meaningfully control AI. I invite the experts to help us better understand the situation and shape the conversations properly.

However, I do think we have to think more carefully about how to contend with the risks AI systems pose when using data that is continually evolving over time. For instance, if an AI model is trained on one fixed dataset and then applied to data that is constantly changing, there is the risk of the model experiencing "drift," which leads to inaccuracy.

There is also the issue that we cannot necessarily regulate what we don't understand. So, if companies bring out a new product with a certain behavior that we might not be able to even understand, how do we think about approving—versus not approving—something as ethical or not ethical?

RAMA CHELLAPPA: The FDA must approve many AI-based procedures. So, what is your opinion about AI regulation? And is it possible?

NATALIA TRAYANOVA: There are several possible overall big-picture approaches. One is that we pump the brakes a bit on the use of deep learning in very high-stakes applications. And the FDA is very much in line with that: They are not very keen to approve black-box deep learning applications for clinical decision-making. The FDA is very interested in explainability in these applications.

For medical applications and beyond, the European Union is taking a more aggressive approach, creating a regulatory framework that sorts potential applications into risk categories. On March 13, 2024, they approved a landmark law governing AI. It imposed a blanket ban on some uses of AI such as biometric-based tools used to ascertain a person's race, sexual orientation, or political beliefs, prohibited the use of deep learning in situations where the potential for harm is high, such as finance, and criminal justice, but allowed low-stakes applications with guardrails.

DANIEL KHASHABI: Like my fellow panelists, I think the bulk of regulation should focus on applications, concentrating on where technology interacts with humans. We have various institutions that regulate and oversee sensitive domains, like the FDA and medical applications and the FAA and aviation.

Given AI's widespread application, we now need another regulatory body with a broad perspective on AI capabilities, deployed across a variety of applications. Maybe this body will work with the FDA and other domain-specific regulators. I worry about overregulating AI development because I don't want our progress to be curbed, but we need regulations on the AI development side, such as a certain level of transparency into how pre-training data was collected, ideally with some level of third-party investigators/researchers.

RAMA CHELLAPPA: I'm with you on that, Daniel. We had a meeting, last week at Johns Hopkins' Bloomberg Center [in Washington, D.C.] with U.S. Senators Mark Warner and Todd Young, and they talked about AI regulation and what Congress wants to do. There are many bills in development, and Senator Warner talked about how AI can negatively impact areas like financial markets. So, he says regulation may be coming, and Senator Young felt we don't have to start something entirely new but could probably build on what is already out there. I agree with a cautious approach to regulating AI.

Do you cover AI in your teaching? Has it changed your approach to teaching?

NATALIA TRAYANOVA: I teach a core undergraduate class that students who are applying to medical school have to take, sort of computational medicine. I include the issues that we talked about today, but it's more general: What does it mean to bring AI into the clinical arena in medicine?

MATHIAS UNBERATH: I teach a class on machine learning and deep learning and one on explainable machine learning, in the context of human-centered design. After [completing] these classes, students should understand the limitations of using these types of approaches.

I built some of my code based on platforms and frameworks that are available online. So, if students want to use generative AI to help them code the homework assignments, given that these large language models are trained on everything that is on the internet, there is a high chance they will do quite well.

But [for quizzes and exams], I've been reverting to the "Middle Ages," where we use paper and a pen, and hopefully, we also use a brain, to demonstrate skills acquired in the class. I have reintroduced closed-book, pen and paper exams.

DANIEL KHASHABI: I always encourage my students in the class or my research lab to use generative AI to increase their productivity. This is just a new tool in their toolbox, to become effective researchers.

This year, we needed to make a slight adaptation to my class. We are careful to make sure that questions we assign for homework are deeper and require more analysis. Even if [students] can prompt [a] language model, they have to be very creative how they prompt the model.

I've also tried to incorporate more interactivity in my class and we do more in-class activities and assignments during class so most learning happens during class hours.

Do you think AI and black-box issues should be taught as part of the first-year curriculum?

NATALIA TRAYANOVA: I don't think in general it's suitable to teach AI and black-box issues in the first-year curriculum as black-box issues in AI require some prior knowledge.

It's important to appropriately gauge when [AI curriculum] would fit a given population of students. You don't want to turn people off or scare them away. You want to introduce knowledge appropriately, so the approach should be to decide on a department-by-department and a school-by-school basis. I think in biomedical engineering it could be taught as part of the undergraduate experience, particularly if ethics issues are emphasized.

MATHIAS UNBERATH: I don't know the right moment to start discussing this. These things move very, very quickly—like when the smartphone was introduced, making this crazy high level of computing accessible at your fingertips.

I don't consider myself old, but when I meet younger people, I realize I'm no longer young. Growing up, we had this modem making funny sounds, and you had to decide whether you wanted to place a telephone call or browse the internet at glacial speeds. We went from that to having a computer in our pockets very quickly.

Professors may be nimble at adapting to challenges and changing their course, but overall, schools might not have moved fast enough in areas such as digital well-being, where we need to consider the implications of disparities in technology use and access.

There are new opportunities and challenges related to how we want to make up the fabric that holds society together: connectedness versus isolation, reverberation of personal views through content, moderation versus engagement with other viewpoints.

So we have to think more broadly, and that might start well before a freshman AI ethics class.

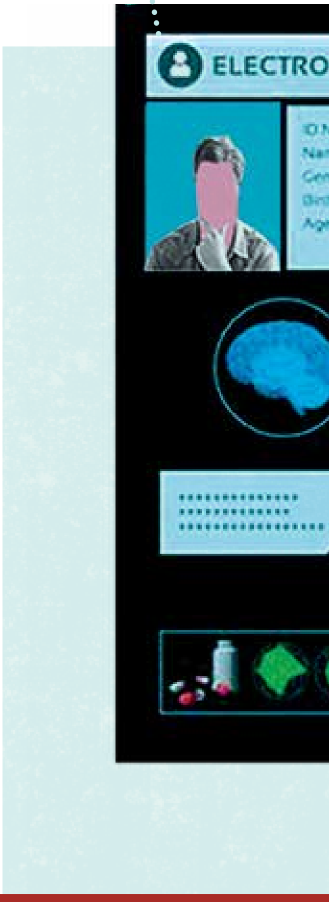

What is a unique or creative use of AI that most people haven't considered yet?

MATHIAS UNBERATH: I would love to see AI help us make more time for personal connections because I think we are spending too much time on tasks that don't require a human touch. If I look at my job or personal life satisfaction, most of it is attributable to interaction with others. Those are usually the most memorable and important moments, but I don't have enough of them because I end up spending too much time on things that I'm not sure I need to be doing. So, this is the one direction where I would love to see AI go.

DANIEL KHASHABI: My first wish is for a personal assistant. We are busy people, and I would love to have an assistant to tell "fix this paper" or do other tasks. That would save me time! My second wish is for trash-collecting robots. Every weekend, I volunteer to collect trash in our neighborhood and collecting trash is going to be a very difficult problem for AI-driven robots since every piece is different. So maybe in a couple of years, I will focus on trash-collecting robots.

NATALIA TRAYANOVA: What I want the most is medical foundation models (generalist medical AI) that make medical decisions without the need to train AI algorithms on specific datasets, then predict specific outcomes, and demonstrate generalizability on external datasets.

Our community has huge difficulties in getting anonymized patient data. So, I would like generalist models developed so we don't need data for each particular medical problem, we don't need to search for months to access external data, and so on. The medical generalist models will be our "base camp"—where we have general medical knowledge, and we just climb from there to solve a specific medical problem [she gestures], not climbing from the bottom every time, developing approaches de novo for each medical problem.

RAMA CHELLAPPA: I think AI will greatly impact engineering design. My hope is to give it a layout of 10,000 buildings and then tell the AI, "Hey, give me a cool new building and these are my constraints. I want it to have four floors. I want this much made of glass," and have it somehow come up with a design that I probably never would have thought of. ■

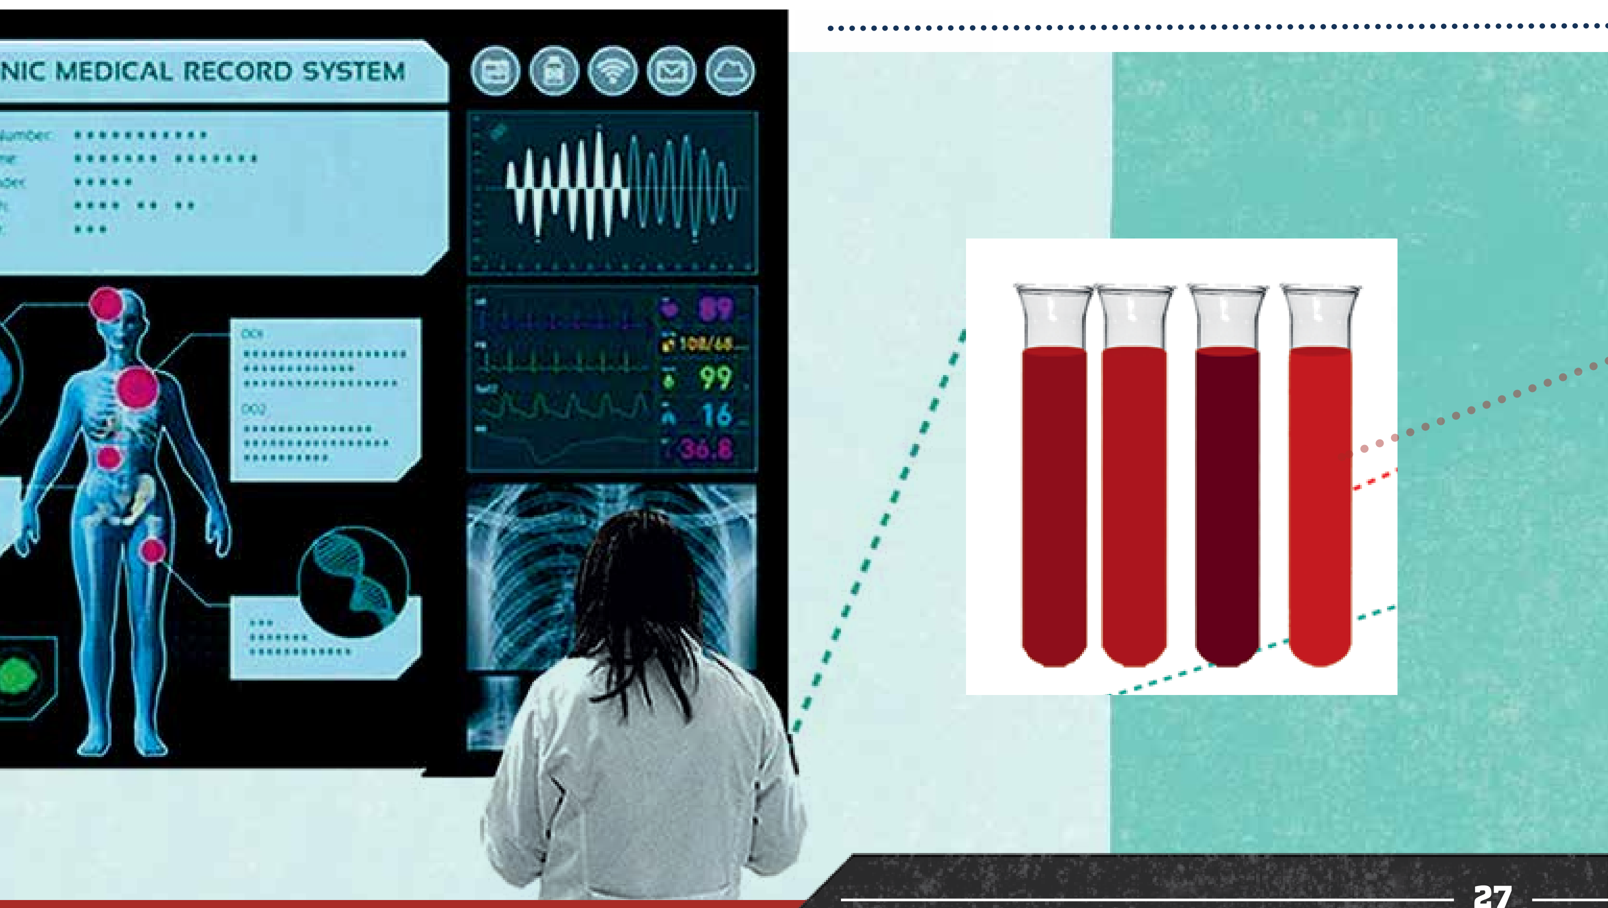

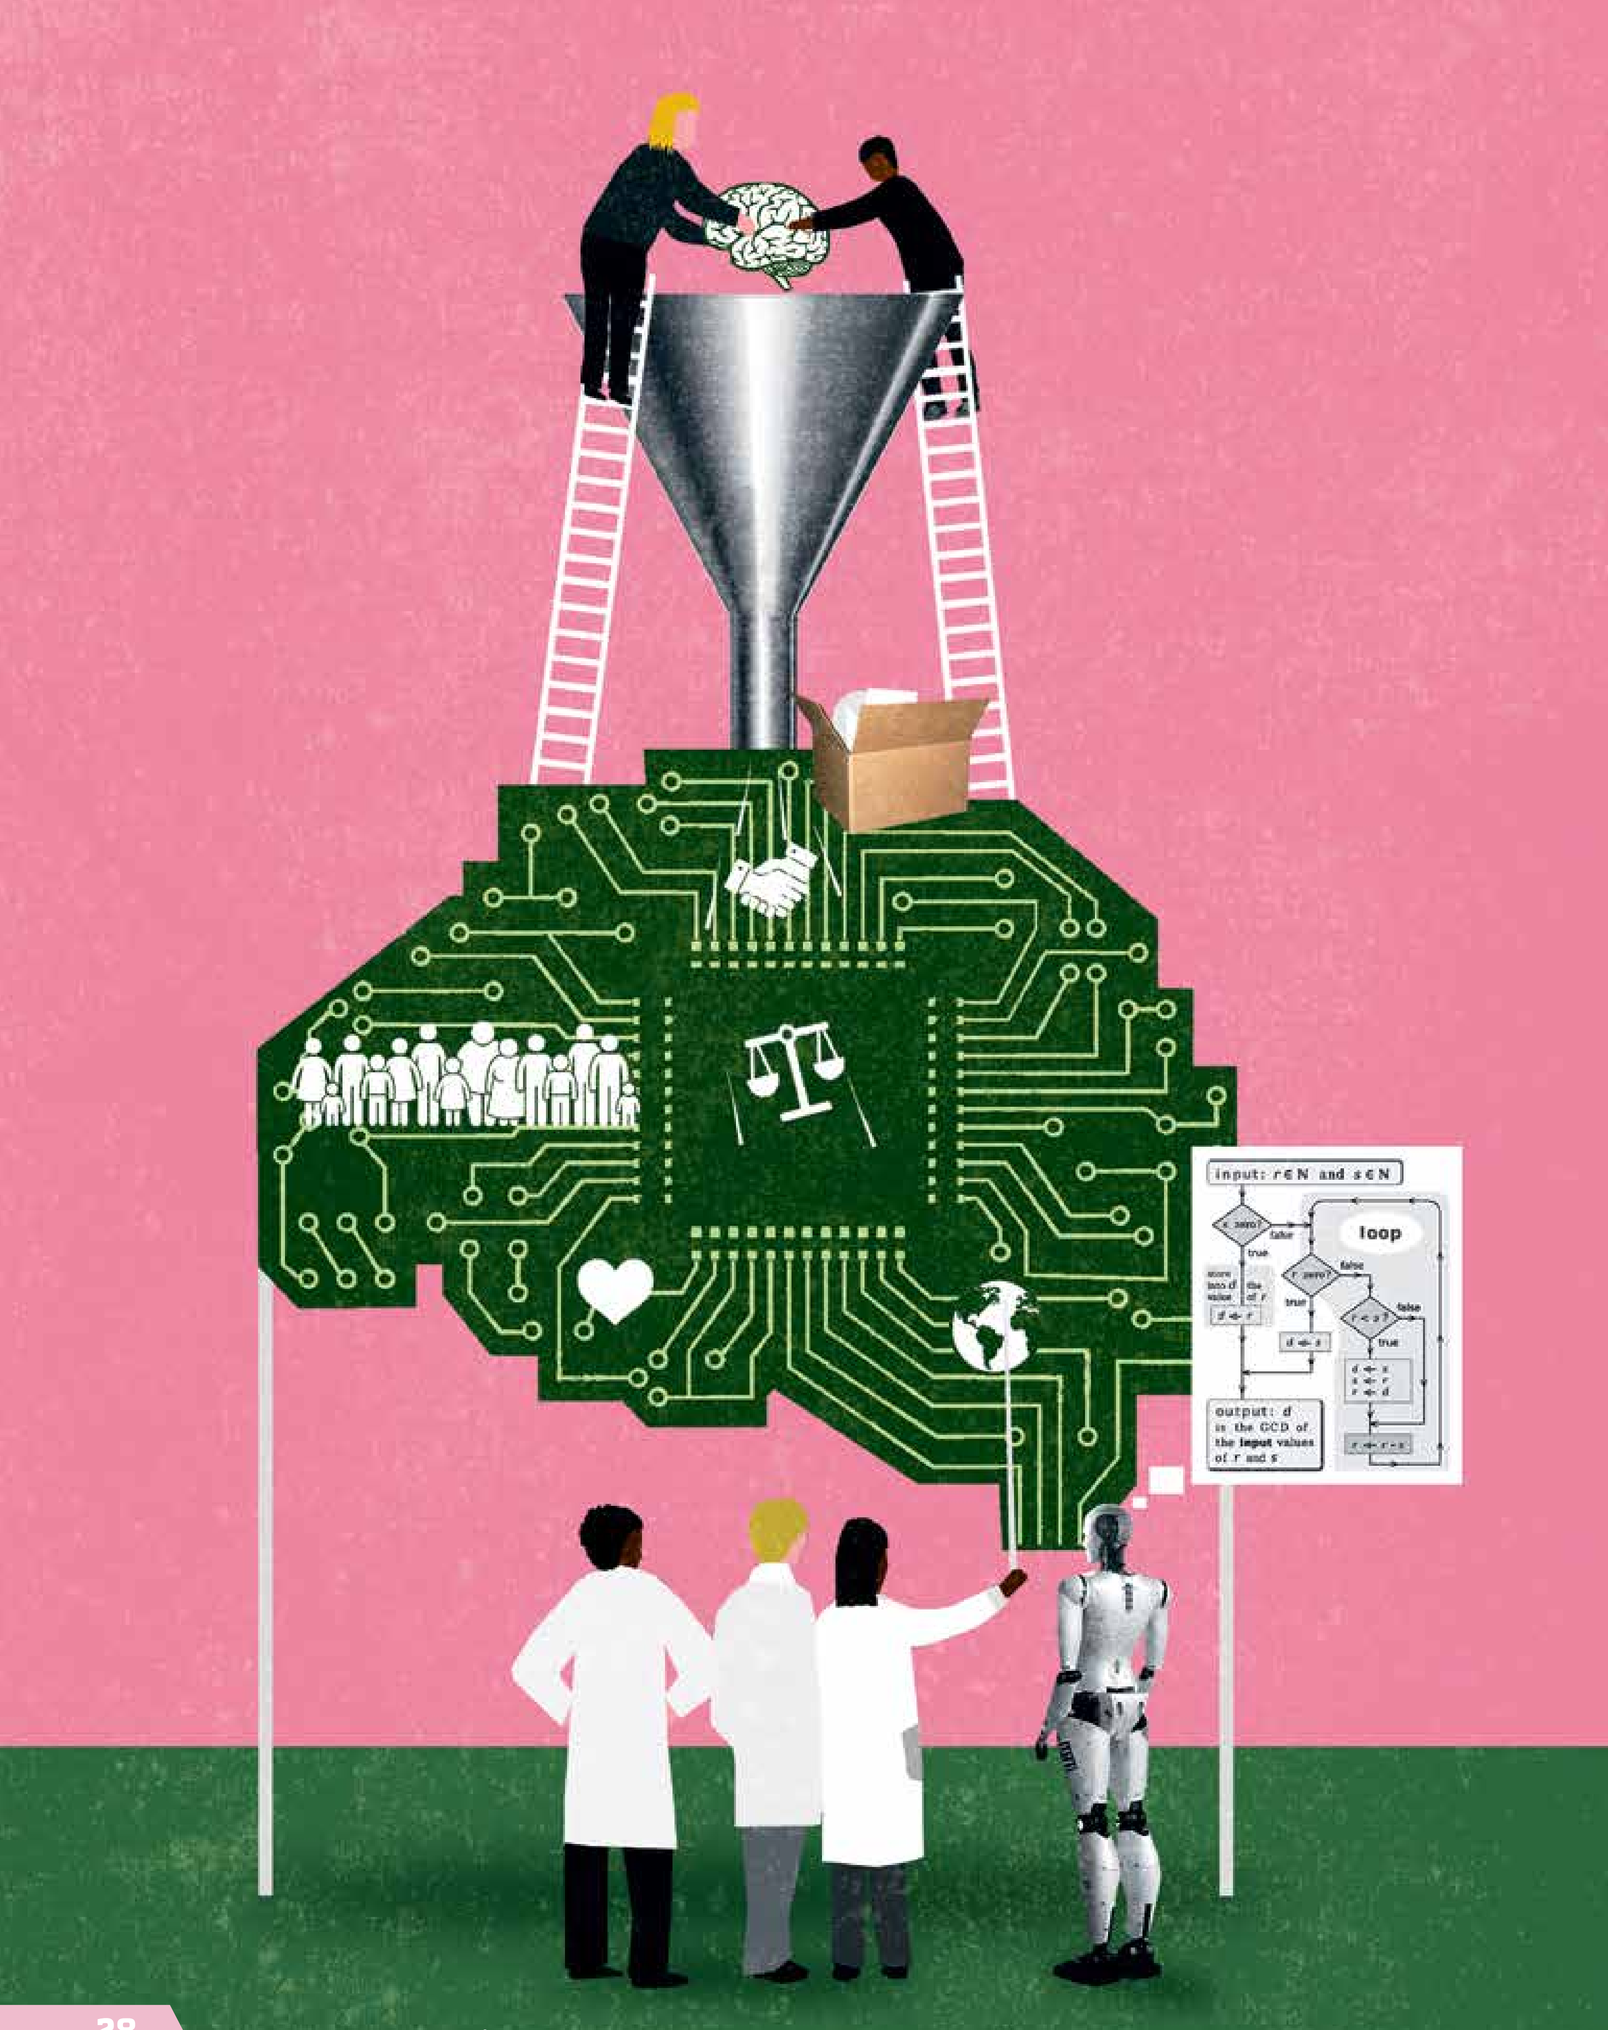
input: r∈N and s∈N
loop
false
true
r zero?
false
true
false
true
output: d
is the GCD of
the input values
of r and s

Delivering on the Promise of PERSONALIZED MEDICNE

BY Adam Hadhazy **ILLUSTRATION BY** Patric Sandri

Harnessing advances in data science and AI, Whiting School researchers are working closely with clinicians to improve care for a broad array of debilitating conditions.

When arriving at diagnoses and treatment plans, physicians often must look past an individual patient's particulars to find the commonalities that are at the root of the disease—in other words, to what makes patients the same. Yet this one-size-fits-all approach can lead to suboptimal care, with drugs and other therapies not working as well as expected for many individuals.

To enhance disease diagnosis and treatment, researchers at the Whiting School of Engineering are instead embracing what makes every patient different. Through advances in data science and AI, reams of patient information can increasingly be gathered by sophisticated yet subtle sensors, both in the clinic and out, and then parsed by algorithms for new insights. In this way, through the development of novel software and hardware, Whiting School researchers are delivering on the promise of precision or personalized medicine—where therapies are tailored to a patient's unique biology.

Here are five ways that data science and AI capabilities are making precision medicine a reality.

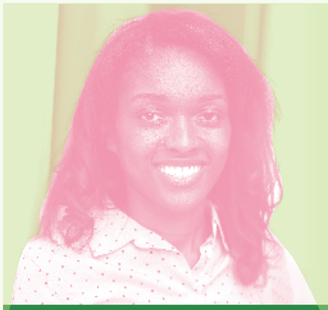

"I went into imaging thinking this is one solution that benefits everyone equally, but that is not the case."

Muyinatu "Bisi" Bell

Making Imaging Equitable

To improve outcomes for patients with cancer, **Muyinatu "Bisi" Bell** figured that proactive detection of the disease in its earliest stages, when it is treatable, was the way to go. She accordingly studied medical imaging, assuming it worked identically for all patients. What Bell quickly learned, though, is that some common imaging methods fail to account for normal human variations in skin tone, tissue type, and tissue density.

"I went into imaging thinking this is one solution that benefits everyone equally, but that is not the case," says Bell, the John C. Malone Associate Professor of Electrical and Computer Engineering, Biomedical Engineering, and Computer Science.

Take ultrasound and photoacoustic imaging, for example. Both these imaging modalities visualize internal anatomy via the detection of sound waves by sensors in contact with the skin. Ultrasound directly bounces sound off internal tissues, while the photoacoustic method shines laser pulses through the skin, and as bodily tissues absorb this laser energy, they subtly expand and generate ultrasound waves.

For both methods, the timing of the receipt of sound wave signals paints a picture of anatomical locations and physiological features, courtesy of data-processing algorithms. Traditionally, these algorithms are premised on a uniform skin tone and standard ratios of tissue components, such as fat to muscle. Yet the skin of darker-complected individuals absorbs greater amounts of light than that of paler individuals, which scrambles photoacoustic imaging results. Likewise, sound moves differently through the tissue layers of patients with varying fat-to-muscle ratios or different densities of breast tissue. The upshot: Pathological features, such as tumors, can be significantly harder to discern and characterize in Black and brown patients, as well as patients with high or low body mass indexes (BMIs) and tissue densities.

Bell is a pioneer in using data science and AI to remove these biases. She devised a new algorithm rooted in the physics of ultrasound wave propagation that instructs sensors to "talk to their neighbors," she says. The sensors compare signal receipt timing and cross-correlate data instead of just lumping signals together based on some impersonalized standard. In this way, Bell's medical imaging algorithm not only constructs a more accurate representation of patients' internal anatomy, regardless of skin tone, BMI, or tissue density, but also offers greater imaging precision for all patients.

One hitch the new algorithm presents is that its complexity requires additional computing resources. As a solution, Bell has leveraged AI to efficiently speed up processing time by training the algorithm on a wide variety of patient data. "For more individualized imaging that is sharper, clearer, accurate, and faithful to ultrasound physics, these more complex algorithms benefit from AI," she says.

Overall, the AI-honed algorithm is inexpensive to deploy and still makes imaging results available in real time—both critical elements from a health care delivery perspective, enabling clinicians in rural and remote areas worldwide to immediately inform patients and arrange appropriate follow-up care.

With continuing development, Bell and her colleagues expect to see imaging systems incorporate their approach in coming years. "We've shown that it is possible to make imaging technology equitable," Bell says.

A Better Way in Epilepsy

Epilepsy, or seizure disorder, ranks among the most common neurological diseases, impacting upward of 60 million people globally. Nearly one in three patients does not respond to drugs, leaving brain surgery as the only mainstream treatment to suppress their seizures. The high-stakes surgery aims to cut out the seizure source, known as the epileptogenic zone (EZ)—a localized area of neural tissue where overactive neurons trigger cascading electrical storms. Yet about half of surgery patients still do not find relief and end up suffering again from seizures just six months post-procedure. This high failure rate stems from how notoriously difficult it is to nail down the exact locations and dimensions of EZs.

"EZTrack can help clinicians achieve the delicate balance of leaving enough of a patient's brain alone so as not to impair the individual, but still taking out enough of the brain to stop the seizures."

Sridevi Sarma

Enter the aptly named EZTrack ("easy track"), a software tool developed by **Sridevi Sarma** and colleagues. EZTrack uses AI—alongside some old-fashioned calculus—to map an epilepsy patient's brain more precisely and uniquely reveal troublesome EZs.

"We've found a way to better identify an epileptogenic zone and carefully understand its boundaries," says Sarma, professor of biomedical engineering and director of the Neuromedical Control Systems Lab. "EZTrack can help clinicians achieve the delicate balance of leaving enough of a patient's brain alone so as not to impair the individual, but still taking out enough of the brain to stop the seizures."

EZTrack capitalizes on voluminous, though presently discarded, data gathered during traditional searches for EZs. For patients with elusive EZs, physicians implant violin-string-like electrical wires, called intracranial electroencephalogram (EEG) sensors, directly into the brain. Patients are then monitored for days to weeks in a hospital setting, where the sensors can trace any bouts of seizure to a distinct brain region.

EZTrack takes a contrary approach by also parsing the massive data haul collected when patients are not seizing. With this information, EZTrack builds a dynamical model of a given patient's brain. That model comprises an interconnected network of nodes, or groups, of neurons. Sarma and colleagues have reasoned that when patients are not experiencing seizures, any fragile, seizure-prone nodes are kept in check by the normal-functioning nodes. With this virtual brain model in hand, EZTrack can simulate stimulation to various nodes, thus zeroing in on fragile nodes as likely EZ candidates. "You can see where a little perturbation caused a huge reaction that never died down, creating a seizure," says Sarma.

Across retrospective studies involving dozens of patients, EZTrack correctly predicted surgical success about 75% of the time—some 25% better than standard methods. Critically, EZTrack can predict surgical failure with near-100% accuracy, meaning the tool could help clinicians and patients avoid procedures that pointlessly remove healthy brain tissue.

An additional advantage of EZTrack, Sarma notes, is that it does not pose any "black-box" problems for patients or clinicians. In AI circles, a black box refers to a program whose inner workings are not readily interpretable by users, or even by the program's developers (read more on p. 23). The issue comes up frequently when AI programs employ machine learning across training sets of data to arrive at efficient, yet inscrutable means of achieving high accuracy. But because EZTrack does not rely on machine learning and is instead a patient-specific computational approach rooted in upper-level math, Sarma says "we can know exactly what the AI is doing."

Having secured a patent for EZTrack along with Food and Drug Administration clearance, the researchers are now working on commercializing the technology. Their AI-powered, brain network mapping could also advance understanding of a range of other conditions including autism, major depressive disorder, and frontal temporal dementia, using data collected by noninvasive EEG sensors placed right on the scalp.

Aging Well

While treating older patients at a Veteran Affairs medical center, **Peter Abadir** realized he had found his calling in clinical geriatrics. "You fall in love with the stories these patients bring, their camaraderie," says Abadir, an associate professor of medicine at Johns Hopkins School of Medicine. "I'm drawn to their history and getting to know the real person behind a name in a chart."

Abadir came to Johns Hopkins in 2007 because it's one of the very few centers in the U.S. that supports both clinical practice and research in geriatrics. In recent years, this intertwining has reached a new level via the Johns Hopkins–based Human Aging Project (HAP). The transdisciplinary initiative is looking toward engineering and biological solutions to lengthen the time that cognitive or functional decline can be kept at bay for older adults.

One of HAP's key pillars is the Artificial Intelligence & Technology Collaboratory for Aging Research (AITC), which Abadir serves as co-principal investigator. Started in November 2021 with $20 million in funding from the National Institute on Aging, the goal of the AITC at Johns Hopkins—along with collaboratories at other hosting institutions worldwide—is to engineer novel tech solutions that boost independent living and reduce physical and cognitive frailty, improving lives of older adults and their caretakers.

Abadir is accordingly helping pool the deep resources on the clinical side of Johns Hopkins with those on the engineering side at the Whiting School. "It's very rare to find such a large number of geriatricians that you can match with engineers," says Abadir. "That's the blessing of being in an ecosystem that has expertise in aging and also engineering."

The day-to-day work of AITC involves funding technology development pilot projects. One core AITC area is Alzheimer's disease and related dementias. An example project in this space is the use of sensing devices in homes and assisted living facilities that can nonintrusively monitor patient activity. That data can power intelligent risk prediction models for common adverse events, such as urinary tract infections and falls, alerting caregivers so they can render earlier, more effective interventions. A second AITC core area is healthy aging. An example project there involves wearable devices to gather mobility, balance, and gait data to screen for patients' physical functioning far more comprehensively than time-limited human physicians could ever do.

"It's very rare to find such a large number of geriatricians that you can match with engineers. That's the blessing of being in an ecosystem that has expertise in aging and also engineering."

Peter Abadir

Overall, Abadir sees "tremendous promise" for AI and data science in geriatrics. Reflecting this potential, the Johns Hopkins University School of Medicine and the Whiting School are planning to jointly open a new 10,000-square-foot hub at the site of a former gymnasium on the Bayview campus in fall 2024, where AITC-affiliated engineers can more readily rub shoulders with clinicians and other HAP groups.

Najim Dehak, an associate professor of electrical and computer engineering and the 2021 Whiting School of Engineering/HAP Scholar, is eager to see the fruits of this synergy with Abadir and colleagues. "The new space offers an exciting opportunity for our engineers to work together with clinicians, nurses, older adults, and their caregivers to come up with technology-driven solutions to some of the biggest challenges older people face, such as social isolation or mobility issues," Dehak says.

The feature continues on page 34.

3 Questions: Filling Gaps, Augmenting Care

Peter Shen ’99, Biomedical Engineering

Interview by Adam Hadhazy

As the North America head of digital health at Siemens Healthineers, a health care solutions company, Peter Shen regularly engages with major health providers and technology partners on issues around adopting innovative technologies in health care. A Johns Hopkins Leadership Fellow for the Whiting School of Engineering, he recently testified before Congress about artificial intelligence in health care.

Tell us about your work at the intersection of health care and technology. How can the latter broadly enhance the former?

It’s the best of both worlds for me—two areas that I’m really passionate about—and now I have the opportunity through Siemens Healthineers to inform doctors and hospitals that “Hey, there's a lot of great technology out there that could help your patients going forward”—everything from data analytics to digital ecosystems and artificial intelligence that we can make practical in the health care arena.

You look at health care as an industry and it’s actually quite behind as it relates to technology. The things we do on a daily basis in our personal lives with technology, like requesting services, banking, communicating—you don’t see that same level in health care. There are good reasons for the slow adoption of technology in health care, primarily because the stakes are higher; we’re talking about patients’ lives, so it’s good to pump the brakes and have guardrails in place. At the same time, it’s an industry that could really benefit from these new and emerging technologies and platforms.

How can AI boost health care?

The industry is strained, partly because we have an aging population that requires more from a health care standpoint. That population’s growth rate compared to the rate of the number of new doctors, nurses, clinical technologists, and other caregivers is not the same. To address this growing workforce gap, we’ve got to figure out a way to be more efficient and maximize the limited time these resources have. AI can help take care of mundane, repetitive tasks so caregivers can do what they do best, which is focus on the patient. It’s a big driver for us as an organization, with the patient at the center of everything we do.

It’s also important to think about using these technologies not just to fill gaps but to augment what’s already there. Patients still trust their doctors, and we want to maintain that trust by helping doctors make better-informed diagnostic decisions and provide more personalized treatments through technology. We have AI solutions delivering quantitative and qualitative clinical information that doctors would not normally have. We need to make sure there’s ubiquitous access to this technology for all patients as well—not just at large health centers in the big cities, but also hospitals, in rural areas, and for underserved populations.

How did you get interested in health care and how did Johns Hopkins nurture that interest?

I grew up in Rochester, Minnesota, a small town that’s known for the Mayo Clinic. My father was an engineer and my mother was a laboratory technologist, so I had an early gravitation toward technology and health care. I heard through friends and relatives: “Hey, there’s this place Johns Hopkins that is kind of good at these types of thing!” [Laughs]

All that led me to my time in Baltimore and I loved every minute of it. I had the opportunity to meet professors and clinicians doing biomedical engineering research and I even got to scrub into the surgical suite during open-heart surgeries. As I reflect on my time at Hopkins, certainly I treasured those experiences, but also now appreciate some of the diverse courses that I took outside of biomedical engineering that really influenced me. I remember my Writing Seminars and history courses, where I learned valuable communication skills and the ability to articulate complex concepts, which has become so important in the role I have today.

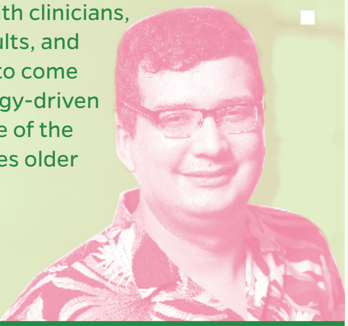

"The new space offers an exciting opportunity for our engineers to work together with clinicians, nurses, older adults, and their caregivers to come up with technology-driven solutions to some of the biggest challenges older people face."

Najim Dehak

Abadir likewise looks forward to seeing the broader endeavor continue to make rapid progress. "AI and other fields are moving at a speed that is unprecedented," says Abadir, "and while that comes with challenges and ethical concerns we'll need to iron out, I think the future will be very different when we'll be using these technologies to enhance our care for older adults."

Deciphering Neurodegenerative Diseases

The emphasis on innovation in geriatric care at Johns Hopkins reflects the skyrocketing rise of this patient demographic. By 2050, the number of people over the age of 60 is forecasted to jump 70% in developed countries—a figure that pales in comparison to an expected quadrupling in less developed regions.

Amid the accompanying spike in diseases of aging, arguably the hardest to adequately treat will be neurodegenerative conditions, which rob patients of their cognitive and physical abilities. "We're likely looking at an epidemic of neurodegenerative diseases," says **Laureano Moro-Velazquez**, an assistant research professor at the Center for Language and Speech Processing. "It's a tough problem to solve and we need new tools."

To this end, Moro-Velazquez is harnessing AI techniques to assist doctors in more efficiently and effectively diagnosing these devastating diseases. Even though Alzheimer's and Parkinson's impact distinctly different neural circuits in the brain, patients can experience similar early symptoms, including trouble communicating, walking, and sleeping. Such overlap complicates precise diagnosis, meaning patients might not get medications matched to their actual pathology, especially in the case of Parkinsonian syndromes.

Moro-Velazquez's work has initially focused on AI-driven analyses of patient speech and language. The models he is developing parse the complexity of a patient's vocabulary and even track pauses—details that can help in disambiguating the underlying neurodegenerative processes. "This level of analysis would take a human hours and hours," says Moro-Velazquez. "With our system, it can take seconds."

In this effort, Moro-Velazquez is working closely with Dehak, and together, the researchers are bringing additional modalities to bear that capture more of the patient's total physical presentation. Besides microphones for speech, high-speed cameras track eye movements as patients read, and other sensors track how they move their hands while writing, and their gait when ambulating.

"We track the whole body, not just speech production," says Dehak. "It's a unique approach that we think can better differentiate one disease from another."

Logging all this information quantitatively further allows for objective comparisons of how patients are doing at follow-up visits even years later—in marked contrast to current practices, where physicians typically consult subjective notes from past patient encounters. Data for AI-enhanced evaluations could also be collected wherever the right hardware is available, better enabling continuity of care should patients [be treated] at different locations, or even conveniently right in their own homes.

"We're likely looking at an epidemic of neurodegenerative diseases. It's a tough problem to solve and we need new tools."

Laureano Moro-Velazquez

Ultimately, the researchers hope their platform will be able to identify telltale biomarkers for neurodegenerative diseases and serve as a protocol for evaluation. "The way you sound, the way you speak, the way you move your eyes—neurodegeneration can affect these things in distinct ways that our models can distinguish, helping doctors establish a diagnosis," says Moro-Velazquez.

Genetic On/Off Switch

Beyond screening and diagnosis, data science and AI may also revolutionize the treatment of disease at the genetic level. The research of **Michael Beer** and colleagues gets to the heart of the matter by investigating enhancers—short regions of DNA that bind certain proteins needed for expressing genes. In this way, enhancers essentially function as genetic on/off switches, playing critical roles in development and health.

"Every cell has the same genetic instruction set, but enhancers allow cells to perform specific functions by reading pieces of that genome out," says Beer, a professor of biomedical engineering at the Whiting School of Engineering.

Despite their importance and ubiquity—hundreds of thousands are predicted to exist within the roughly three billion "letters" of the human genome—enhancers are elusive entities. Unlike other genomic features, enhancers do not seemingly have any explicit markers. Making matters worse, the specialized snippets do not necessarily appear next to the genes they regulate.

To combat this challenge of enhancer identification, Beer has been developing machine learning programs over his nearly 20-year career at Johns Hopkins. Continued honing of algorithms and experimental methods has recently led to significant advances. In a paper published in *Nature Genetics* last summer, Beer and colleagues focused on the brief but intense periods of activity when cells transition from one cell type to another. Such bursts prominently occur during embryonic development for cell fate decisions, when undifferentiated stem cells transition into distinct cell types.

In their study, Beer's group examined the rise of endodermal cells, which go on to specialize further as gut-lining cells and other cellular populations. As a first step, the researchers constructed a network model of gene regulation. They then used a technique based on the gene-editing tool called CRISPR to block the activity of suspected enhancers in stem cells cultured in the lab. Next, the researchers monitored for any effects on subsequent gene expression as the stem cells morphed into endodermal cells.

With a resulting wealth of data in hand, the researchers modeled the physical state of the DNA to

"Every cell has the same genetic instruction set, but enhancers allow cells to perform specific functions by reading pieces of that genome out."

Michael Beer

look for anything that might better reveal furtive enhancers. One big takeaway: Even though enhancers can lie distantly upstream or downstream in DNA from genes they act upon, physical looping of DNA back in on itself usually places enhancers and genetic correspondents in close proximity. "We saw that enhancers were clustered inside DNA loops," says Beer. Detecting and digging down into these loops could therefore prove fruitful in enhancer discovery.

A second, even bigger takeaway: Many enhancers strongly influence gene expression only during the brief transition phase, but not afterwards during the bulk of a cell's life. This finding helps solve a key mystery about how previous genome-wide studies have highlighted thousands of disease-associated enhancers that, when tested, show perplexingly little effect on gene expression. "Our results explain a lot about why enhancers are so hard to study," says Beer. "Their effect is only at the right time."

Thanks to Beer's computational approach, geneticists can hope to track down which enhancers potently contribute to disease, and from there pursue enhancer-targeted therapies to control gene expression. Diabetes and cancer are just two of the conditions that could become more tractable if dealt with via fundamental genetics—especially at the personalized granularity of enhancer errors. "I think cancer is a disease of genetic regulation," says Beer. "Genetic mutations are clearly needed, but what mutations are doing at some level is regulated by the state of the cell. And the state of cell is controlled by enhancers."

Glow and Sew could enable surgeons to see sutures through blood and tissue.

No Ordinary Suture

JUSTIN OPFERMANN, A MECHANICAL ENGINEERING PHD STUDENT IN THE IMERSE Lab, watches a monitor in Hackerman Hall's mock operating room. On the screen, he observes as his colleague **Yaning Wang** carefully stitches two pieces of synthetic tissue together using the da Vinci robotic surgical system.

Wang guides the two robotic forceps as they pass a clear suture back and forth, threading it through the pink rubber. But with a couple of taps on the monitor to change the camera settings, Opfermann demonstrates that this is no ordinary suture. Everything on the screen turns black and white except for the stitches, which glow bright green.

The fluorescent suture is part of Glow and Sew, a tool Opfermann and his partners Wang and **Jiawei Ge** are developing to help surgeons improve outcomes in a surgical procedure known as anastomosis that involves joining two tubes together. This surgery can be challenging, especially when there's blood in the surgical field or the tissue folds in on itself, making the suture difficult to see and leading to less-than-perfect sewing.

"It's very difficult, but you need to be perfect at it 100% of the time to stop a complication," such as sepsis, from occurring, Opfermann explains.

Unlike typical sutures, the Glow and Sew suture is coated in a fluorescent powder that glows under near-infrared light, which is also able to penetrate through the blood and tissue, allowing surgeons to always see where their stitches are. "It essentially is giving you a superhuman capability to see something that's normally hidden in the surgical process," Opfermann says.

Opfermann came up with the idea for Glow and Sew with the help of **Axel Krieger**, PhD '08, an associate professor of mechanical engineering, and **Simon Leonard**, assistant research professor in computer science, around five years ago. But it wasn't until earlier this year that he was able to test his idea, thanks to funding from the Whiting School's Student Initiatives Fund. Glow and Sew received $2,500 in support to buy supplies, develop and test initial concepts, and create a working prototype that the team can use to secure additional funding.

Next, Opfermann aims to find a way to incorporate artificial intelligence and machine learning to create overlays that show surgeons where each stitch should go. He plans to use the data from this research to apply for patents and other grants with the goal of bringing his ideas into actual operating rooms to improve patient outcomes.

—ABIGAIL SUSSMAN

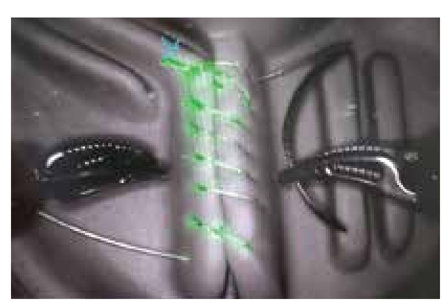

TELL US WHAT YOU THINK!

Continue the conversation online, using **#JHUEngineering** on Facebook and Twitter!

Advocating for Students With Disabilities


Ikshu Pandey

IKSHU PANDEY'S INTERESTS LIE AT THE INTERSECTION OF MATERIALS SCIENCE AND neuroscience.

"The big reason that I am interested in these two majors is because I'm disabled," says Pandey. "Because of a virus that I caught in my childhood, I can't really see out of my left eye or hear from my right ear."

As a senior double major, Pandey strives to assist other students with disabilities and neurodegenerative conditions in both her academic work and through her extracurricular activities—efforts recognized with a 2023 JHU Diversity Leadership Award that honors outstanding work in diversity and inclusion at Johns Hopkins.

For her senior materials science and engineering senior design project, Pandey worked in the lab of **Peter Searson**, the Joseph R. and Lynn C. Reynolds Professor in materials science and engineering and the co-founder and former co-director of the Institute for NanoBioTechnology, on research that aims to provide earlier detection of Alzheimer's disease. She has collaborated on designing microfluidic devices that mimic blood vessels in the brain. These biomaterial molds can be injected with very small amounts of biological fluid, allowing researchers to compare changes in blood vessel models with Alzheimer's disease to those without the condition.

"I'm seeding [the devices] with stem cells that have familial mutations associated with Alzheimer's disease and characterizing the changes that occur," says Pandey.

As the undergraduate representative on the inaugural Johns Hopkins Disability Inclusion Advisory Committee, she advocates on behalf of other students. "I'm representing my undergraduate population and making sure that students with disabilities aren't overlooked," she says.

Pandey is also a student director for Baltimore First, a grassroots organization dedicated to building relationships between Hopkins students and members of the Baltimore community.

"Through engaging with the Baltimore community, I realized that within health care, there are a lot of stigmas," Pandey says, referring to race-based stigmas and how they sometimes intersect with stigmas attached to disabilities. "If doctors are not able to engage sensitively with a patient in a way that makes [the patient] comfortable enough to tell them their whole story, [the doctor] might misdiagnose them."

And, as the director of community engagement for a new, youth-led national nonprofit, Asian Americans with Disabilities Initiative (AADI), she helped compile the Asian Americans with Disabilities Resource Guide.

Pandley plans to attend medical school and believes that by working in medicine, she can help build a future where equitable access to all in engineering and medicine is afforded. "Science doesn't only rely on time-tested procedures and ideas. It also relies just as much on different perspectives and new ideas," she says. "Through this experience, Pandey grew more aware of the range of stigmas that exist in the practice of health care, including race-based stigmas and those related to disabilities."

— CONNER ALLEN

Greg Wulffen

> "The first thing we did when tackling this project was to think of all the ways a chimp might try to destroy it, and focus in on that."
> —GREG WULFFEN

A Primate-Proof Monkey Bar

IN THEORY, TAKING AN EKG JUST REQUIRES PLACING SMALL PLASTIC-TIPPED ELECTRODES ON a patient's skin to detect electrical impulses from the heart. But the procedure becomes far more complex when the patients are active, curious chimpanzees—who suffer from similar heart conditions to humans.

The challenge facing **Greg Wulffen**, a combined BS in mechanical engineering and MS in engineering management student: Design an EKG tool that could be used on chimps at the Maryland Zoo—or, rather, used by the chimps themselves. The zoo's one requirement: Make it primate-proof.

"The first thing we did when tackling this project was to think of all the ways a chimp might try to destroy it, and focus in on that," says Wulffen, who inherited the Multidisciplinary Engineering Design project—known as PulseApe—from a team that had started it three years earlier.

Alissa Burkholder Murphy, senior lecturer in the Center for Leadership Education and director of the Multidisciplinary Design program, revived the project last year. She says that Wulffen and **Mirza Korman** '23, took the original team's ingenious prototype to the next level: "They hit it out of the park."

At issue were the delicate electronics housed at the center of the bar, which were not yet well-protected "from the fingers or teeth of the prying primates." Using Computer-Aided Design (CAD) software, Wulffen came up with a sheath made of Delrin—a thermoplastic that has a high-fracture toughness—that slides over a cut-out in the center for the electronics and coated the handles with copper. The team then made the entire device wireless, using Bluetooth, and refined the interface so a zookeeper could easily decode and extract the EKG data.

After the zoo confirmed that all materials were safe, the time came to try out the device on the intended subject: an energetic chimpanzee.

"The first run-through overall went well, and we got a reading from one of the chimps," says Pam Carter, area manager for the Chimpanzee Forest at the Maryland Zoo. "I think we're the only place getting EKGs this way."

Wulffen is now using his expertise for a Multidisciplinary Design project at the National Aquarium, where he is part of a team designing a durable, corrosion-resistant device to allow visitors to interact with a new underwater Wetlands Exhibit. This summer, he will take those problem-solving skills to Tesla, where he is interning with the company's Technical Program Management team.

—JONATHAN DEUTSCHMAN

What will your legacy be?

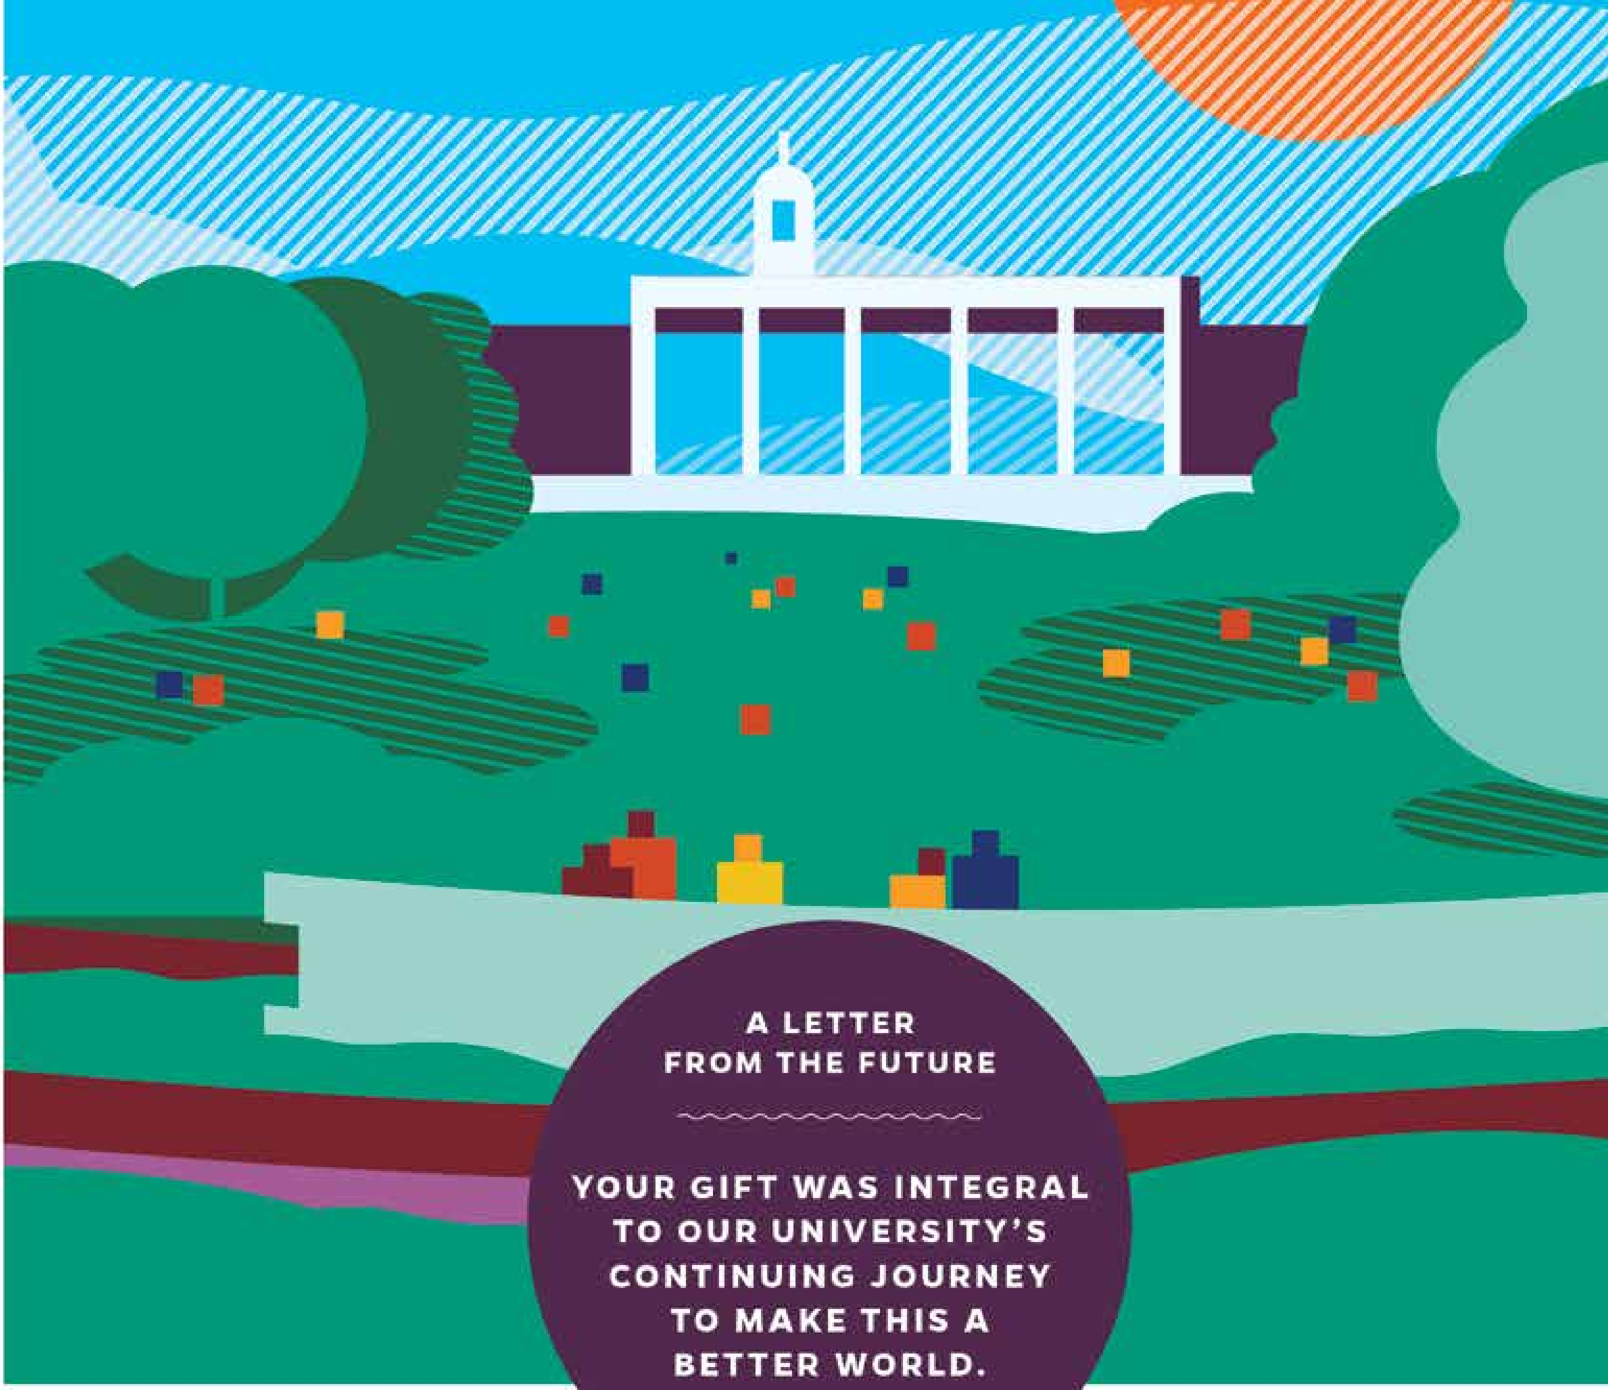

A LETTER FROM THE FUTURE

YOUR GIFT WAS INTEGRAL TO OUR UNIVERSITY'S CONTINUING JOURNEY TO MAKE THIS A BETTER WORLD.

THANKS!

A gift of retirement assets is an easy and efficient way to give back to the Whiting School. Plus, you and your family could benefit from significant tax advantages.

JOHNS HOPKINS
WHITING SCHOOL *of* ENGINEERING

To learn more, contact the Office of Gift Planning:
410-516-7954
or 800-548-1268

giftplanning@jhu.edu
giving.jhu.edu/giftplanning

Jeong Kim receiving the National Medal of Technology and Innovation from President Joe Biden, Fall 2023.

Kim Honored With National Medal of Technology

LAST FALL, JEONG KIM '82, MS '89, RECEIVED THE NATIONAL MEDAL OF TECHNOLOGY AND Innovation, the country's highest award for technological achievement.

The award recognized Kim's contributions to data communications and broadband optical systems, including the invention of the asynchronous transfer mode (ATM) switch for wireless communications, a groundbreaking innovation that allowed the transfer of data, voice, and video over a single network infrastructure.

Just two decades after arriving as a 14-year-old immigrant from Korea with limited English, Kim sold his first company, Yurie Systems, for $1 billion in 1998. Today, he remains a leading innovator in telecommunications technology.

Kim credits education for much of his success. With the help of a Whiting School scholarship, he earned bachelor's degrees in electrical engineering and computer science, followed by a master's degree in technical management at Johns Hopkins. "Hopkins is academically strong, and they gave me the flexibility to take as many credits as I wanted so I could finish in three years," he says.

He next served as a nuclear submarine officer in the U.S. Navy for seven years, obtaining his master's degree while on shore duty at the Defense Nuclear Agency, and later earned a PhD in reliability engineering from the University of Maryland. It was after the Navy, while working as a contract engineer at the Naval Research Laboratory, that Kim developed the asynchronous transfer technology that enabled multimedia battlefield communications, the innovation that laid the groundwork for the ATM switch.

Today, Kim is co-founder and chairman of Kiswe Mobile, a live-streaming and interactive video platform for sporting events and concerts. Prior to that role, he served eight years as the 11th president of Bell Labs and was the first leader to ever be recruited from outside the organization. He was also jointly appointed as a professor in the mechanical engineering and the electrical and computer engineering departments at the University of Maryland and has served as a Johns Hopkins trustee and member of Whiting School's Advisory Board.

"Being recognized by the President of the United States for accomplishments was not my intention," says Kim, referring to the National Medal of Technology and Innovation. "But to have the things that I've done actually have made an impact on the lives of other people is very rewarding personally. I feel honored that I was able to make that impact."

— ROBIN SMITH

THIS PAGE: PHOTO COURTESY OF RYAN K. MORRIS AND THE NATIONAL SCIENCE & TECHNOLOGY MEDALS FOUNDATION. OPPOSITE PAGE: COURTESY OF PEDIAMETRIX.

TELL US WHAT YOU THINK!

Continue the conversation online, using **#JHUEngineering** on Facebook and Twitter!

A Better Measure for Infants

Interview by Robin Smith

Fereshteh Aalamifar

In 2018, Fereshteh Aalamifar, who earned her PhD in electrical and computer engineering in 2016, and her husband Reza Seifabadi launched SoftSpot, now the only FDA-cleared mobile app for infant cranial measurements. SoftSpot assists pediatricians and parents with early detection, monitoring, and treatment of cranial conditions, which—if left untreated—can affect babies' brain development in some cases.

WHAT COMPELLED YOU TO TACKLE THIS PARTICULAR HEALTH CARE CHALLENGE?

My passion for children, my own struggles as a new mother, and my desire to make health care accessible to families inspired me to found PediaMetrix. After I had my first child, I learned that a significant amount of skull growth happens in the first months and how crucial it is during that time to look at the skull's size and shape as it grows.

HOW DOES YOUR APPROACH TO SOLVING THIS PROBLEM DIFFER FROM OTHER METHODS?

Pediatricians relied on visual assessments and according to some statistics, up to 600,000 infants miss the window of early detection of cranial deformities every year in the U.S. SoftSpot runs on almost all mobile phone devices, enabling parents and clinicians to upload photos of a baby's head and receive a report about the child's skull shape.

SoftSpot enables users to make informed decisions earlier and more easily because it can continuously provide objective and quantitative data when sporadic, subjective, or visual assessment is the only alternative. However, it does not replace a clinical professional.

WHAT OBSTACLES DID YOU ENCOUNTER WHILE DEVELOPING THIS TECHNOLOGY?

Getting access to clinical data and learning and then implementing regulatory compliance were the biggest challenges in the early years. Thanks to advocates, we were able to collaborate with several children's hospitals in the U.S. to conduct clinical studies.

WHAT OTHER AREAS OF HEALTH CARE DO you **ENVISION YOUR TECHNOLOGY BEING APPLIED TO?**

Our next version of SoftSpot will use a 3D model of a baby's head to identify and plan for cases that require surgery and to eliminate the need for additional scans that can expose a baby to radiation.

WHAT ARE THE NEXT FRONTIERS OR CHALLENGES YOU HOPE TO ADDRESS IN THE HEALTH CARE SPACE?

My passion lies in using technology advances to create equal access to health and wellness for people in all communities, including low-resource populations. Above all, I would like to contribute to ensuring every child gets the best start in life.

World Changer

Quinton Smith

IN DECEMBER 2023, QUINTON SMITH, PHD '17, WAS HERALDED AS ONE OF THE YEAR'S "TOP 10 SCIENTISTS on the cusp of changing the world" by *Popular Science* magazine. The periodical's annual "Brilliant 10" honors early-career researchers under the age of 40 who stand out as innovators and changemakers in their fields.

Smith works at the vanguard of engineering and stem cell biology research in efforts to create "organoids"—tissues grown from stem cells that mimic our organs and can be used to show how organs interact with drugs.

The science behind personalized regenerative medicine begins when "you take a mature cell from various origins such as the skin, hair, or fat, and reprogram it into an embryonic state," says Smith, who joined the Henry Samueli School of Engineering at the University of California, Irvine in 2021 as an assistant professor in chemical and biomolecular engineering. "This so-called human induced pluripotent stem cell (hiPSC) technology offers the ability to regenerate any cell in the human body."

By adding certain chemical factors, Smith and his team of six PhD and two postdoctoral scholars in his Integrated Developmental Design Laboratory can change hiPSCs into any type of human cell. "We're [currently] making a representation of a liver—a so-called 'organ on a chip'—without the ethical constraints that come with human embryonic stem cell research," says Smith. He notes that the liver has more than 500 different functions, including the metabolization of drugs. "By creating liver tissue from stem cells," he says, "we can study drug efficacy to develop new medications for treating ailments, such as nonalcoholic fatty liver disease."

While earning his undergraduate degree in chemical engineering from the University of New Mexico, Smith attended the Johns Hopkins Institute for NanoBioTechnology (INBT) Research Experience for Undergraduate Program. He says that his passion for stem cell research began during the two summers he spent under the guidance of Sharon Gerecht, a former Johns Hopkins faculty member and director of INBT, who later became his PhD advisor when he joined Johns Hopkins for his graduate degree.

Smith says the recent honor from *Popular Science* is a humbling experience. "I'm just a product of the great mentorship and research environments that I've been fortunate to be involved in."

— NIKOLAS CHARLES

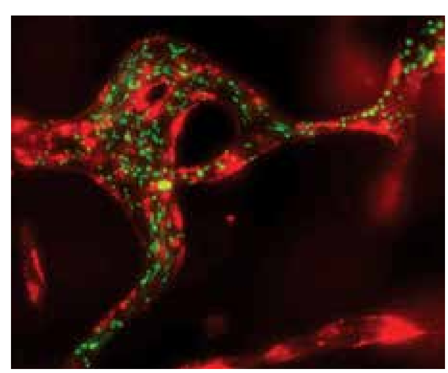

Beads through engineered vessel

THIS PAGE: PHOTO COURTESY OF QUINTONSMITH; OPPOSITE PAGE: PHOTO COURTESY OF JESSICA JEFFERY.

Clearing the Air

Jessica Jeffery

BY THE TIME JESSICA JEFFERY, MS '19, ENTERED MIDDLE SCHOOL IN BUTTE, Montana, she had already decided to pursue a career in the field of environmental engineering. Living in an area with polluted water and land due to a century of copper mining had a strong impact on her.

"I wanted to make sure that others didn't have to grow up with mining waste right out their back door," says Jeffery, who earned her master's degree in environmental planning and management from the Whiting School's Engineering for Professionals (EP) program.

Today, as a staff environmental engineer with General Motors in Detroit, Jeffery works cross-functionally with executives, regulatory agencies, and corporate partners as she provides technical direction and regulatory guidance to GM manufacturing facilities across the country.

One of her recent projects was overseeing the construction of Ultium Cells, a battery cell manufacturing facility in Warren, Ohio, in a joint venture between GM and LG Energy Solution. "About 90% of my work is focused on electric vehicles and battery cell facilities," says Jeffery, who is currently providing her environmental regulatory expertise for another battery cell facility in Spring Hill, Tennessee.

As a subject matter expert for environmental regulations related to the construction and operation of new battery cell manufacturing facilities, Jeffery's areas of expertise include the in-depth understanding of and compliance with all U.S. Environmental Protection Agency acts, including the Clean Air Act and the Toxic Substances Control Act, a law that ensures new chemicals are safe for use.

Jeffery says the EP's part-time curriculum allowed her to pursue her advanced degree while working full time. Out of several standout online classes, Kelly Tzoumis' Environmental Justice was one of the most noteworthy, as the instructor taught in a "unique and impactful" style. Jeffery said that Tzoumis used various media to highlight areas of the country suffering a disproportionate amount of pollution and the laws in place to ensure equal protection for those communities.

"I continue to learn new skills in my manufacturing career," says Jeffery. "And I still find it fascinating that you can take a piece of metal and make an entire vehicle out of it. It's also a privilege to work with my fellow GM employees and the teams working at the manufacturing facilities. My goal is to continue to find ways to strive for clean air while leveraging technology to improve our impact on the environment."

— NC

Good Vibrations

Howard Katz

> "There are many connections between musical instruments and materials science."
> —HOWARD KATZ

A PIONEER IN THE FIELD OF ORGANIC ELECTRONICS AND PHOTONICS, HOWARD KATZ OCCASIONALLY turns to his trusty cello to provide his engineering students with an object lesson in wave vibration.

"There are many connections between musical instruments and materials science," says Katz, a materials science and engineering professor, who most recently performed Robert Schumann's cello concerto for students in his Electronic, Optical, and Magnetic Properties of Materials class.

"Many instruments incorporate carefully designed materials to generate the sound and amplify it," he explains. "For example, the bow hair has a very complicated structure of tiny fibers that grab the strings and cause them to vibrate."

Katz's perfect pitch—the ability to recognize and imitate exact musical notes—sparked an initial interest in the piano.

"I cannot remember a time that I could not play the piano by ear…. Around age 8, I began taking piano lessons," he recalls. "The school I attended in West Orange, New Jersey, had one piano, and many students wanted to be piano players in the orchestra. I decided to pick an orchestra instrument so I could always be in the orchestra."

He soon began studying the cello, igniting a passion for performance that continued into adulthood. While earning degrees in chemistry and music theory from MIT, Katz was the lead cellist for the MIT Symphony Orchestra, playing in trios and quartets throughout college. He still finds opportunities to play, both inside and outside the classroom.

In June, Katz often joins the Baltimore Symphony Academy, a weeklong program that brings amateur artists together with professional musicians in the Baltimore Symphony Orchestra for side-by-side rehearsals and performances.

To Katz, participation in the academy is equivalent to playing in the Major League. "For some people it's a once-in-a-lifetime dream to do this, to say that they had a Major League Baseball experience," he says. "I'm very lucky, being in Baltimore. I get to do it every year."

— CONNER ALLEN

THIS PAGEAND BACK COVER: WILL KIRK/HOMEWOODPHOTO.JHU.EDU.


Synbotics
GRP 2400
APL

JOHNS HOPKINS
WHITING SCHOOL
of ENGINEERING

3400 N. CHARLES ST.
SAN MARTIN CENTER, 1ST FLOOR
BALTIMORE, MD 21218

Acoustic Aviary

In spring, Johns Hopkins' Homewood campus is a symphony of chirping birds, students chattering, and the hum of lawnmowers and hedge trimmers. To a group of engineering students, this pleasant auditory chorus is more than background noise: It's an opportunity to use acoustic sensors to learn more about avian life on campus.

The electrical and computer engineering undergraduates installed a network of devices in trees on campus. These sensors use radio links to transmit data about sound sources, rather than actual audio, according to **Andreas Andreou**, professor of electrical and computer engineering and project mentor.

"What we are doing in this Homewood Soundscape project is monitoring birds and learning about the types of birds on campus: basically, deciphering their chirps and calls and turning it into meaningful data."

The project combines sensory systems, embedded processing, and machine learning: all crucial electrical and computer engineering concepts. Eventually, the students plan to expand the project to include other sensors, including those that measure water quality in the surrounding streams.

"Ultimately, this marriage of Internet of Things (IoT) technology and environmental monitoring showcases JHU's dedication to harnessing the power of technology for the betterment of its community and beyond," Andreou said.

— **LISA ERCOLANO**