EQuad News
Winter 2019
Volume 30, Number 2

Supplement to the *Princeton Alumni Weekly*

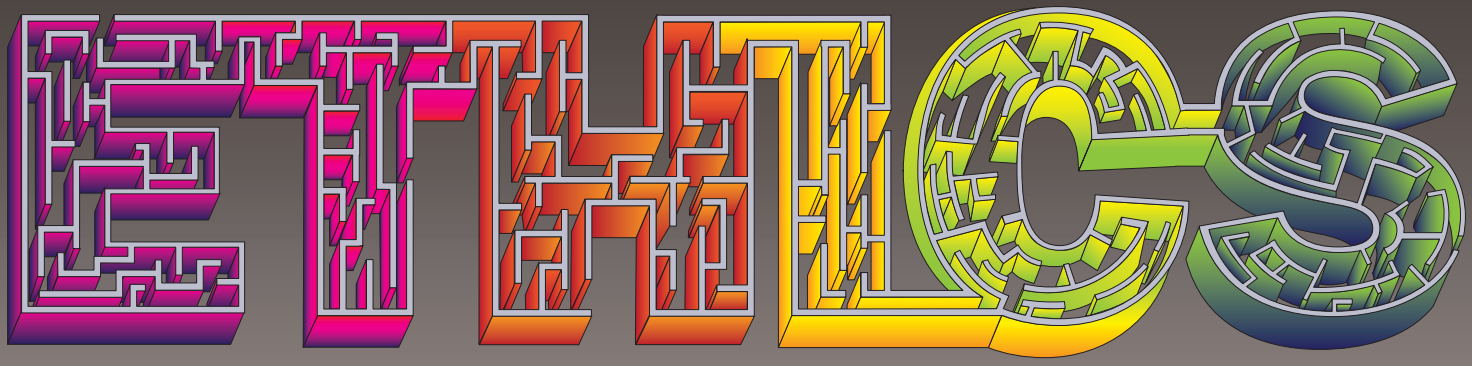

in Engineering

PRINCETON
School of Engineering and Applied Science

Ethics in Engineering

Making ethical choices, crucial in any walk of life, is especially important in engineering. The work of engineers is grounded in fundamental science but translates quickly into structures, devices, and systems that bring broad benefits and big changes to the world. Most recently, the sweep of digital technologies across our lives has raised questions around privacy and security, even while increasing productivity and convenience. Similar issues arise in any engineering discipline. More than a century of development around fuels, transportation, buildings, and medicine — to name a few areas — has remade many aspects of daily life, the environment, the global climate, and human health, while promising myriad future advances. With almost every product, bridge, or vehicle we use, we trust in the ethical behavior of engineers to keep us safe.

When the stakes are this large, the demand for ethical behavior and integrity is even greater.

In my own research group, I make ethics an explicit topic of conversation with students and postdocs, not only as we go about day-to-day research and publishing but as we discuss the motivation behind our work and the implications of our findings. We discuss conflicts of interest and encounter the fuzzy lines and cultural differences that emerge when the precision of math and science intersects with the variations of human nature. I am proud to lead an engineering school where these conversations are common. And I am grateful to be part of a university where faculty and students collaborate and learn easily across disciplines, including public policy and the humanities, which are needed for a view of technology framed from a human perspective.

This magazine offers a sampling of ways in which ethics is an important part of research and teaching in the School of Engineering and Applied Science.

Where have ethics and technology intersected in your professional lives? Write to us at eqn@princeton.edu or talk with us via Twitter, Facebook, or Instagram.

Emily Carter
Dean
Gerhard R. Andlinger Professor in Energy and the Environment
Professor of Mechanical and Aerospace Engineering and Applied and Computational Mathematics

Photo of Emily Carter by David Kelly Crow

Contents

950096

EQuad News
Winter 2019
Volume 30, Number 2

Dean
Emily Carter

Vice Dean
Antoine Kahn *78

Associate Dean, Undergraduate Affairs
Peter Bogucki

Associate Dean, Development
Jane Maggard

Associate Dean, Diversity and Inclusion
Julie Yun

Director for Administration, Finance, and Planning
Kaitlin Lutz

Director of Engineering Communications
Steven Schultz

Senior Editor
John Sullivan

Web Designer
Neil Adelantar

Digital Media Editor
Aaron Nathans

Writer
Molly Sharlach

Communications Specialist
Scott Lyon

Contributors
James Bronzan
Liz Fuller-Wright
Adam Hadhazy
Molly Seltzer

Additional Editing
Liz Fuller-Wright

Graphic Designer
Matilda Luk

EQuad News is published twice a year by the Office of Engineering Communications in collaboration with the Princeton University Office of Communications.

Note on alumni class years
Following Princeton University convention, undergraduate alumni are indicated by an apostrophe and class year; graduate alumni, whether master's or doctoral, are indicated with a star and class year.

EQuad News
C-222, EQuad,
Princeton University
Princeton, NJ 08544

T 609 258 4597
F 609 258 6744
eqn@princeton.edu

engineering.princeton.edu/news/e-quad-news

In the Nation's Service and the Service of Humanity

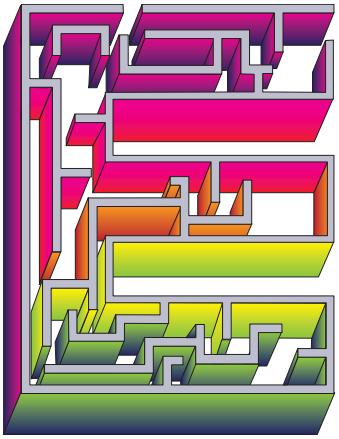

PRINCETON'S FIRST AI4ALL SUMMER PROGRAM AIMS TO DIVERSIFY THE FIELD OF ARTIFICIAL INTELLIGENCE

For three weeks last summer, high school students packed the conference rooms of Princeton's computer science building, honing programming skills while taking on challenges in artificial intelligence — from sharpening computer vision for self-driving cars to identifying fake or misleading online news.

Princeton's inaugural AI4ALL summer program brought 32 rising 11th graders to campus for an immersive residential camp, July 23 to Aug. 11. The program brings young people from underrepresented groups into the growing field of artificial intelligence (AI), a branch of computer science that involves creating autonomous systems capable of processing human language, recognizing images, or analyzing complex data sets.

The students learned the fundamentals of AI technology through lessons and hands-on exercises, and heard from Princeton professors working at the forefront of AI research in computer vision, cybersecurity, and bioinformatics. They also delved into their own group research projects, with guidance from Princeton graduate students and recent graduates. The projects explored a variety of AI applications: smart speakers and home appliances — including the gadgets' security vulnerabilities; algorithms that guide self-driving cars; processing human language to flag false news; and making meaning of massive data sets from the Fragile Families and Child Wellbeing Study, which has tracked outcomes for nearly 5,000 children from birth to age 15.

"We really try to emphasize AI for social good," said program co-director Olga Russakovsky, an assistant professor of computer science. Russakovsky established the AI4ALL foundation in 2017 along with her former Ph.D. adviser, Fei-Fei Li, a 1999 Princeton alumna and a professor of computer science at Stanford University.

Last year, AI4ALL expanded to six universities in North America, with each camp featuring a different focus. Princeton's program has been developed as a joint effort of the AI4ALL foundation, the Department of Computer Science, and the Center for Information Technology Policy (CITP). The program is additionally sponsored by the Siegel Family Endowment with contributions from the Center for Statistics and Machine Learning and the Office of Information Technology.

To help diversify the next generation of AI researchers and developers, Russakovsky and co-director Ed Felten, director of CITP, recruited students from groups underrepresented in AI who performed well in high-level math and science courses at their schools. "But beyond that," said Russakovsky, "we really look for students who are going to be leaders in any field that they choose to pursue."

– by Molly Sharlach

Princeton's AI4ALL summer program aims to bring young people from underrepresented groups into the growing field of artificial intelligence (AI). During a three-week residential camp, rising 11th graders explored AI applications such as smart speakers and home appliances — including the gadgets' security vulnerabilities. (Photo by Nick Donnoli)

CHIP RAMPS UP ARTIFICIAL INTELLIGENCE SYSTEMS' PERFORMANCE

Princeton researchers, collaborating with Analog Devices, have fabricated a chip that markedly boosts the performance and efficiency of neural networks — computer algorithms modeled on the workings of the human brain.

In tests, the chip performed tens to hundreds of times better than other advanced, neural-network chips.

The researchers believe that the chip could help advance image recognition and other neural-network applications, including artificial intelligence systems in vehicles and robots.

"This kind of improved performance could let mobile devices do intensive tasks, like recognizing their owner's face, without taking up too much time or eating into the device's battery life," said lead author Hossein Valavi, a graduate student in the lab of co-author Naveen Verma, an associate professor of electrical engineering.

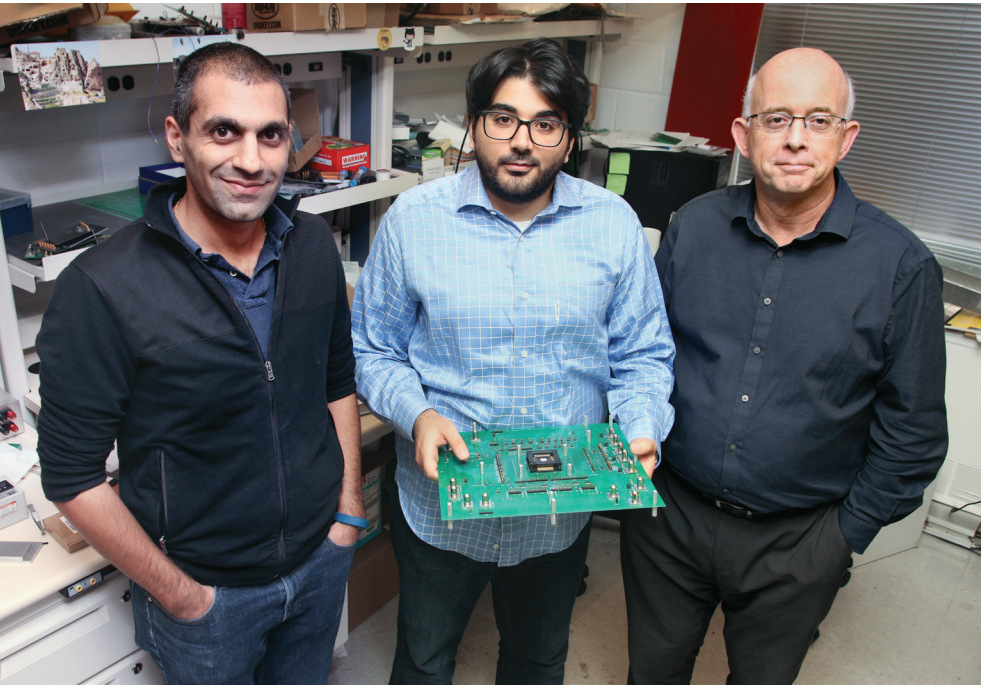

From left: Naveen Verma, Hossein Valavi, and Peter Ramadge have fabricated a chip that markedly boosts the performance and efficiency of neural networks — computer algorithms modeled on the workings of the human brain. (Photo by Frank Wojciechowski)

Other authors of the study, which was published in IEEE Symposium on VLSI Circuits, are Peter Ramadge, the Gordon Y.S. Wu Professor of Engineering and director of the Center for Statistics and Machine Learning, and Eric Nestler of Analog Devices, a Massachusetts-based semiconductor company.

Artificial neural networks are complexes of connected units that can be trained to make decisions from data given in structurally complex forms. A key component of neural-network systems is accelerator chips, which boost computational performance. But the accelerator chips themselves can suffer from bottlenecking due to the heavy data flows coursing through their components.

The researchers' chip works with a technique, called in-memory computing, which saves energy and time by performing computations where data is stored, rather than moving it to a different location.

Verma's team followed up the project with a new in-memory chip that is fully programmable. They believe it could be particularly useful on phones, watches, or other devices that have limited battery life but rely on high-performance computing.

Verma said that for many applications, the chip's energy savings would be as critical as the performance boost. But both improvements are only of use if they can be accessed by the broad base of applications that need them — that is where the need for programmability comes in. The chip now works with common computer languages such as C.

"The previous chip was a strong and powerful engine," said Hongyang Jia, a graduate student in Verma's group and one of the designers. "This chip is the whole car."

– by Adam Hadhazy

CLASSIC MATH PROBLEM PROVIDES NEW INSIGHTS FOR FIRST STEPS OF LIFE

For centuries, scientists have puzzled over configurations of objects that fit a given space. These packing problems not only have fascinated mathematicians, they also hold critical implications for disciplines from chemistry to shipping.

Now, researchers at Princeton and MIT have applied lessons from classical packing problems to provide insights into the early steps of animal development. In an article in the journal Nature Physics, the researchers describe how they used advanced microscopy and mathematical modeling to examine the 3D organization of cells in a fruit fly egg chamber — the multicellular precursor to the fertilizable egg. The researchers found that certain packing configurations are favored in the egg chamber, and that these patterns can be classified mathematically.

The findings have important implications for developmental biology because the physical arrangement of the cells within the egg chamber can play a critical role in the development of the egg and, eventually, the embryo.

For their study, the researchers examined roughly 160 egg chambers from fruit flies. Each chamber is composed of a cluster of 16 cells packed in an enclosure of tissue. The cells are connected by bridges of cellular material called ring canals, and the connections follow the same pattern in all egg chambers. If the connections were laid out into a two-dimensional grid, they would form a highly reproducible — what scientists call stereotypic — branching pattern.

The researchers used computer models to examine this tree-like formation in three dimensions to determine whether there are recognizable patterns in the configuration. Jasmin Imran Alsous *18, one of the lead authors, said the fruit fly chambers were particularly well suited to this experiment.

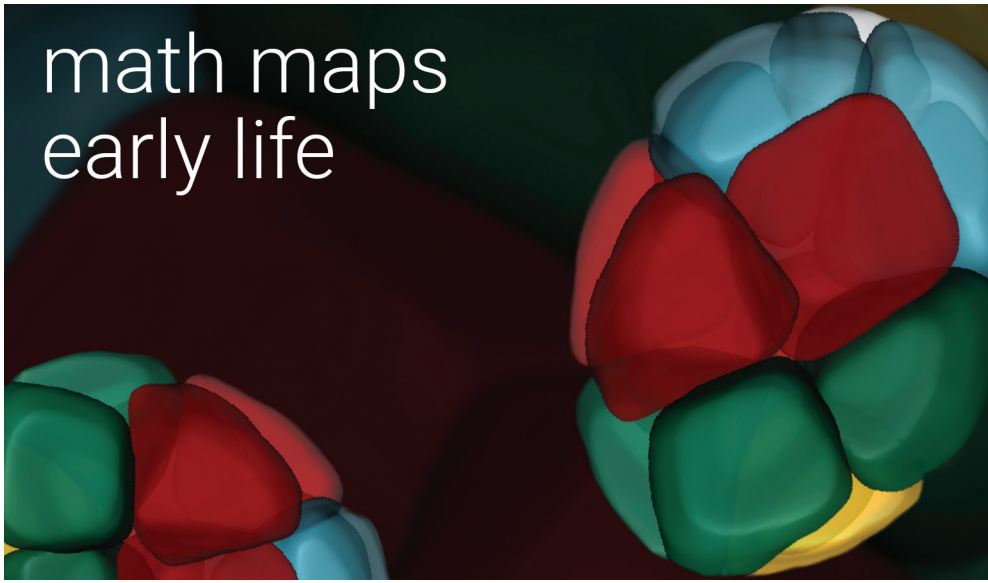

(Illustration by Princeton University)

"This is a beautiful system that is remarkably tractable, where you can identify each cell and study how it is positioned within the cluster," she said.

The math in the analysis is complex, but the researchers concluded that some packing arrangements were more likely than others. Specifically, higher entropy configurations were observed more commonly than lower ones.

"All germ cells, eggs and sperm, in organisms develop in small clusters, like the one we analyzed in this study," said Stanislav Shvartsman *99, a professor of chemical and biological engineering and the Lewis-Sigler Institute for Integrative Genomics and a leader of the research team. "The more we understand about them, the closer we get to some of the first steps in organismal development."

– by Scott Lyon

‘FOCUSED RESEARCH TEAMS’ TAKE ON EMERGING OPPORTUNITIES IN BIOTECHNOLOGY AND ROBOTICS

One of the engineering school’s new Focused Research Teams will look for new generations of antibiotics, in part by studying protective chemicals that populations of helpful bacteria, such as the ones pictured here, produce to ward off harmful invaders. (Image courtesy of the Link Lab)

In an initiative to boost collaborations on subjects too new to fit into existing departments and centers, the School of Engineering and Applied Science has created a program to fund small, cross-disciplinary groups of researchers called Focused Research Teams.

The school has named three initial teams, two in emerging areas of biotechnology and one in robotics and “cyber-physical” systems. Each will receive $250,000 per year for three years, after which they will be evaluated to determine whether the initiative should continue, evolve into a larger effort, or conclude.

“The pace of discovery and level of creativity among our faculty is astonishing,” said Emily Carter, dean of engineering. “And much of this work happens when people from different disciplines begin working together and inspiring each other. In our recent strategic planning process, we identified the need to incubate and accelerate the most exciting new areas so we can more quickly bring their benefits to society.”

Precision Antibiotics: Developing a new generation of antibiotics that avoids drug resistance while protecting helpful bacteria that are often killed by conventional antibiotics. The work, which includes faculty from chemical and biological engineering (CBE) and molecular biology (MolBio), is supported by a fund established by Helen Shipley Hunt *71.

Engineering Living Organelles: Understanding how components within cells develop and engineering them to treat disease or perform new functions such as producing biofuels. The work, which includes faculty from CBE, mechanical and aerospace engineering (MAE), and MolBio, is supported by a fund established by Lydia and William Addy ’82.

Robotics and Cyber-Physical Systems: Advancing development of robotic systems that work alongside humans or in distributed, interconnected groups, starting with a model fleet of robots that collaborate to collect trash. The work, which includes faculty from computer science, electrical engineering, and MAE, is supported by the Addy Fund.

– by Steven Schultz

FELTEN CONFIRMED AS MEMBER OF U.S. PRIVACY AND CIVIL LIBERTIES OVERSIGHT BOARD

Ed Felten (Photo by David Kelly Crow)

The U.S. Senate confirmed on Oct. 11 the appointment of Ed Felten to the Privacy and Civil Liberties Oversight Board, a bipartisan agency within the executive branch.

Felten, the Robert E. Kahn Professor of Computer Science and Public Affairs and director of the Center for Information Technology Policy, will serve as a part-time member of the board while continuing his teaching and research at Princeton. The five-person board, which was established at the recommendation of the 9/11 Commission, is charged with evaluating and advising the executive branch on anti-terrorism measures with respect to privacy and civil liberties.

The appointment adds to Felten’s extensive experience in the federal government. Felten served in 2011 and 2012 as chief technologist at the U.S. Federal Trade Commission and from 2015 to 2017 as deputy chief technology officer in the Obama White House, where he worked on national security issues and led work on the implications for society caused by the growing capabilities of artificial intelligence. **– by Steven Schultz**

NOBEL LAUREATE FRANCES ARNOLD TALKS WITH STUDENTS AND RESEARCHERS ON CAMPUS

Shortly after Frances Arnold '79 won the Nobel Prize in Chemistry, the mechanical and aerospace engineering graduate spent a day on campus talking with students, postdocs, and faculty.

"Your ability to move and adapt [gives] you the flexibility and also the creativity that allows you to put crazy new ideas together," she told an engineering school gathering. "So I hope you will take the opportunity while you are at this magnificent learning institution to fill your thoughts with things outside of what you are doing in the laboratory and really do learn to read, write, and think broadly."

Arnold, a professor at the California Institute of Technology, is the first Princeton alumna to win a Nobel and the fifth woman to win the chemistry prize. She was recognized "for the directed evolution of enzymes." The proteins produced in her lab can speed up reactions that don't exist in nature, leading to more environmentally friendly manufacturing of chemical substances — including medicines — and renewable fuels for a greener transport sector. **– by Liz Fuller-Wright and Steven Schultz**

Frances Arnold (left) visits campus. (Photo by Steven Schultz)

DAY OF OPTIMIZATION EXPLORES FOREFRONT OF DATA SCIENCE

How should society decide who gets a liver transplant? Should there be marketplaces for data, and how should these markets run? If a driverless car kills someone, who is at fault?

More than 350 researchers and students considered these questions at the inaugural Princeton Day of Optimization on Sept. 28. Speakers covered topics including machine learning, methods for maximizing efficiencies of systems, and theories underlying computer control of systems.

"The day turned out to be a success beyond our most optimistic projections," said Amir Ali Ahmadi, the organizer and an assistant professor of operations research and financial engineering (ORFE).

The conference topics are increasingly prevalent in people's lives. A good example is the driverless car. To succeed, the cars need to be efficient, well controlled, and able to deal with a vast array of situations.

"Technology, and to some extent society, will increasingly be relying on advanced algorithms, computation, and data to ensure available resources are used most efficiently," said Peter Ramadge, director of Princeton's Center for Statistics and Machine Learning and the Gordon Y.S. Wu Professor of Engineering. Ronnie Sircar, department chair of ORFE, said planning is underway for the 2020 Day of Optimization, "but we are going to need a larger auditorium." **– by James Bronzan**

Student volunteers check in participants. (Photo by Frank Wojciechowski)

Researchers using satellite imaging have found much greater than expected deforestation since 2000 in the highlands of Southeast Asia, a critically important world ecosystem. (Photo courtesy of the researchers)

SOUTHEAST ASIAN FOREST LOSS MUCH GREATER THAN EXPECTED, WITH NEGATIVE IMPLICATIONS FOR CLIMATE

Satellite imaging has revealed that deforestation in the highlands of Southeast Asia, a critically important world ecosystem, has proceeded much more quickly than expected since 2000. The findings are important because they raise questions about key assumptions made in projections of global climate change as well as concerns about environmental conditions in Southeast Asia in the future.

Zhenzhong Zeng, a postdoctoral research associate in civil and environmental engineering and the lead author of an article describing the findings in Nature Geoscience, said the researchers used a combination of satellite data and computational algorithms to reach their conclusions. The report shows a loss of 29.3 million hectares of forest (roughly 113,000 square miles or about twice the size of New York State) between 2000 and 2014. Zeng said that represents 57 percent more loss than current estimations of deforestation made by the Intergovernmental Panel on Climate Change. He said most of the forest has been cleared for crops.

Crops grow in a recently deforested area in Southeast Asia. (Photo courtesy of the researchers)

Because forests absorb atmospheric carbon, and burning forests contribute carbon to the atmosphere, loss of forests could be devastating. An accurate estimation of forest cover is also critical for assessments of climate change.

Transformation of mountainous regions from old forest to cropland can also have widespread environmental impacts from soil retention to water quality in the region.

Eric Wood, the Susan Dod Brown Professor of Civil and Environmental Engineering and a senior member of the research team, said the results were troubling in that farmers are carving new agricultural frontiers from the highland forests of mainland Southeast Asia. “These forests are an important source for sequestering carbon as well as critical water sources for adjacent lowlands,” he said.

– by John Sullivan

FROM 'SEA OF MUTATIONS,' POSSIBLE CANCER LINKS RISE TO THE SURFACE

By analyzing data from thousands of patients, Princeton researchers have identified genetic mutations that frequently occur in people with uterine cancer, colorectal cancer, or skin cancer — an important step toward using genome sequences to better understand cancer and guide new treatments.

Many genetic malfunctions are commonly found in cancers, and scientists are working to distinguish those directly involved in the initial growth of cancer and its spread. The complexity of both the disease and the ordinary functions of the genome have made this a difficult problem.

"There are numerous mutations that occur in any individual's cancer," said Mona Singh, senior author of the paper and a professor of computer science and the Lewis-Sigler Institute for Integrative Genomics (LSI). "A major question in cancer genomics is to figure out, from this large sea of mutations, which ones are actually causal for cancer initiation, but also for progression."

In a study published in the journal PLOS Computational Biology, Singh's group addressed this question by lining up mutations in a large group of genes that play key roles in gene expression regulation and other cellular processes. These genes are known as zinc finger genes because they have a protein structure that involves a zinc ion. The researchers examined mutations in these genes using the Cancer Genome Atlas, which includes more than 10,000 tumor samples from 32 cancer types.

Daniel Munro, a graduate student in LSI and the paper's lead author, said the researchers used genes' similarities to find common cancer mutations that disrupt protein structures, leading to impaired cellular functions.

Cancer cells have high mutation rates, but the researchers found that certain sites in zinc finger genes were mutated at even higher rates than this background level. Taken together, the results are strong evidence that "zinc finger genes are major players in the gene dysregulation you see in cancers," said Singh.

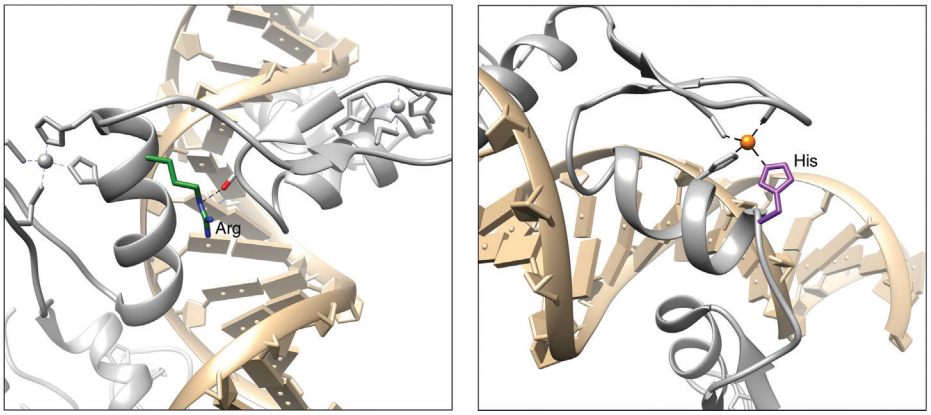

By analyzing thousands of patients' data, Princeton researchers led by Mona Singh have identified two genetic mutations occurring at high frequencies in three types of cancer. The mutations affect amino acids in protein molecules shown in the panels above. (Image by the researchers)

Singh said she hopes the results will help guide laboratory research and lead to progress in understanding the mechanisms of cancer. "We're keen to have other people follow up on this work," she said.

The research was supported in part by the National Science Foundation, the National Institutes of Health, the State of Nebraska, and the Nebraska Research Initiative.

– by Molly Sharlach

Mona Singh (Photo by Frank Wojciechowski)

TIMING IS KEY FOR BACTERIA SURVIVING ANTIBIOTICS

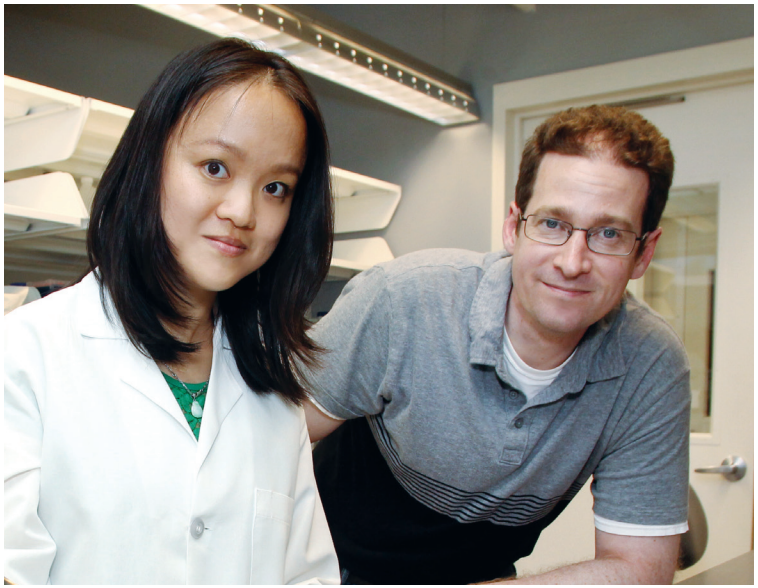

Wendy Mok and Mark Brynildsen are studying how bacteria survive antibiotics, with the goal of making antibiotic treatment more effective. (Photo by Frank Wojciechowski)

For bacteria facing a dose of antibiotics, timing might be the key to evading destruction. In a series of experiments, Princeton researchers found that cells that repaired DNA damaged by antibiotics before resuming growth had a much better chance of surviving treatment.

When antibiotics hit a population of bacteria, often a small fraction of "persister" cells survive to pose a threat of recurrent infection. Unlike bacteria with genetic resistance to antibiotics, evidence suggests that persisters stay alive in part by stalling cellular processes targeted by the drugs.

The researchers examined antibiotics that target bacterial DNA. In bacterial populations, some cells repair damaged DNA before resuming growth, and others resume growth before making repairs. The researchers found that those that make repairs before resuming growth generally are the ones that survive as persisters. The research advances a long-term goal to make antibiotic treatment more effective, and is now part of a new Focused Research Team developing "precision antibiotics" (see story on page 4).

In results published in the Proceedings of the National Academy of Sciences, Wendy Mok, a postdoctoral researcher, and Mark Brynildsen, an associate professor of chemical and biological engineering, analyzed the responses of *E. coli* bacteria to treatment with the antibiotic ofloxacin. Their work built on previous results from Brynildsen's lab that revealed that persisters to ofloxacin required DNA repair machinery to survive.

"But that doesn't guarantee that they're necessarily going to survive," said Mok. "We hypothesized that the timing of DNA repair and the resumption of growth-related activities like DNA synthesis could impact the survival of persisters after treatment."

Mok and Brynildsen used a strain of *E. coli* genetically engineered to allow researchers to control the cells' growth. The researchers used the bacteria to create a uniform population of cells with stalled growth that could tolerate the ofloxacin antibiotic.

These non-growing cells, they found, experienced DNA damage similar to growing cells treated with ofloxacin. However, the non-growing cells showed delays in resuming DNA synthesis and repair following treatment.

Mok and Brynildsen examined persistence in normal cells placed in a low-nutrient environment to stall their growth, simulating a condition that bacteria frequently encounter within an infected host. Indeed, following ofloxacin treatment, if cells were starved of carbon sources for at least three hours, they observed nearly complete tolerance to the antibiotic. This tolerance depended on effective DNA repair processes.

"Nutrient starvation is a stress that bacteria can routinely encounter at an infection site," said Mok. "Our results suggest that in the period after antibiotic treatment we can consider looking at targeting some of these DNA repair processes, and see whether that can improve treatment outcomes." **– by Molly Sharlach**

EFFORT PUSHES COMPUTER SCIENTISTS TO CONSIDER USES OF TECHNOLOGY

by John Sullivan

Computer courses usually concentrate on technology, but Princeton researchers are revamping parts of the computer science curriculum to teach students how the technologies they develop will affect society.

"We are teaching people how to design and implement technology that we will not only use but also depend on," said Nick Feamster, a professor of computer science who is experimenting with adding several ethics modules to the undergraduate computer networking course. "To merely teach the technology, without compelling students to ask questions about the potential ethical implications of the technology, is a problem."

Feamster added ethics instruction to his networking class in 2017, and he is working with others in the computer science department to add similar material to other courses. The effort emerged from the senior thesis of Jasmine Peled, who graduated with a degree in computer science in 2018. She began the project after taking several courses from Feamster.

"The thesis at first looked at how other people had incorporated ethics into their classes," said Peled, who now works on computer networks for the U.S. Department of Defense. "We decided to have the thesis focus more specifically on best practices."

Peled looked at a variety of materials available for teaching ethics to computer scientists. She also interviewed students about ethics and tested teaching methods to see which were most effective.

"We ran six different workshops to test them," Peled said.

Peled found that in general, students preferred to learn about ethics from computer scientists rather than specialists brought in from outside departments such as philosophy or law. She said the students also responded to a variety of presentation styles, including case studies, lectures, and readings.

Feamster said he has seen growing interest in ethics from computer scientists — professors, students, and practitioners alike. In addition to simply creating new technology, leaders of the field also have to think about the application of the technology. He said the ethics instruction would begin in networking and computer security classes with an eye to expanding to other courses if possible. Feamster and Peled are also seeking to create a template that can be used beyond Princeton.

"We want the material to be modular and scalable," Feamster said. "We would like to create material that others at other universities can also use in their computer science classes."

Nick Feamster and Jasmine Peled '18 are working to add instruction about ethics to the computer science curriculum at Princeton. (Photo by David Kelly Crow)

PRINCETON COLLABORATION BRINGS NEW INSIGHTS TO THE ETHICS OF ARTIFICIAL INTELLIGENCE

by Molly Sharlach

Edward Felten (left), director of the Center for Information Technology Policy, and politics professor Melissa Lane (right), director of the University Center for Human Values, created the "Princeton Dialogues on AI and Ethics." (Photo by Sameer A. Khan/ Fotobuddy)

Should machines decide who gets a heart transplant? Or how long a person will stay in prison?

The growing use of artificial intelligence in both everyday life and life-altering decisions brings up complex questions of fairness, privacy, and accountability. Surrendering human authority to machines raises concerns for many people. At the same time, AI technologies have the potential to help society move beyond human biases and make better use of limited resources.

"Princeton Dialogues on AI and Ethics" is an interdisciplinary research project that addresses these issues, bringing engineers and policymakers into conversation with ethicists, philosophers, and other scholars. At the project's first workshop in fall 2017, watching these experts get together and share ideas was "like nothing I'd seen before," said Ed Felten, director of Princeton's Center for Information Technology Policy (CITP). "There was a vision for what this collaboration could be that really locked into place."

The project is a joint venture of CITP and the University Center for Human Values, which serves as "a forum that convenes scholars across the University to address questions of ethics and value" in diverse settings, said director Melissa Lane, the Class of 1943 Professor of Politics. Efforts have included a public conference, held in March 2018, as well as more specialized workshops beginning in 2017 that have convened experts to develop case studies, consider questions related to criminal justice, and draw lessons from the study of bioethics.

"Our vision is to take ethics seriously as a discipline, as a body of knowledge, and to try to take advantage of what humanity has understood over millennia of thinking about ethics, and apply it to emerging technologies," said Felten, Princeton's Robert E. Kahn Professor of Computer Science and Public Affairs. He emphasized that the careful implementation of AI systems can be an opportunity "to achieve better outcomes with less bias and less risk. It's important not to see this as an entirely negative situation."

Ethical knowledge is particularly critical for decisions about AI technologies, which can influence people's lives at a greater speed and on a larger scale than many previous innovations, but risk doing so with insufficient accountability, said Lane. Felten cited the use of automated risk assessment or prediction tools in the criminal justice system to make decisions about bail, sentencing, or parole — decisions that are traditionally made by human judges.

One major question is whether AI systems should be designed to reproduce the current human decision patterns, even where those decision patterns are known to be profoundly affected by various biases and injustices, or should seek to achieve a greater degree of fairness. But then, what is fairness?

Philosophers have always known that there are different ways to view fairness, said Lane. "But we haven't always been pressed to work out the implications of committing to one view and operationalizing it in the form of an algorithm. How do we evaluate those choices in a real-life setting?"

The project's initial case studies, released in May 2018 and available for public use under a Creative Commons license, are based on real-world situations that have been fictionalized for study purposes. They examine ethical dilemmas arising from various applications of AI technology. CITP visiting research fellows Chloé Bakalar and Bendert Zevenbergen played key roles in coordinating and authoring the case studies, with input from the project's workshop participants.

Lane noted that the case studies are intended as a starting point for conversations on AI and ethics in classroom settings, as well as among practitioners and policymakers. Beyond these specific examples, she said, "we also are very conscious of the society-wide, systemic questions — the questions about monopoly power, the questions about privacy, the questions about governmental regulation."

The project will seek to broaden the scope of its resources and teaching tools in the coming years, with the goal of building a new field of research and practice. Collaborations with similar efforts by other universities are also underway, including joint conferences with Stanford and Harvard held in fall 2018.

Case Studies

More information about the project and the case studies is available at **https://aiethics.princeton.edu**

Should a high school collect data on its students' behaviors to identify at-risk teens? What is the appropriate balance between privacy and improving educational outcomes, and who should decide?

If a sound-recognition app correctly identifies details, including a speaker's gender, 99.84 percent of the time, would the error rate be acceptable if used at such a large scale that numerous people are embarrassingly misidentified each day?

A chatbot helps law enforcement officials catch online identity thieves, but is the bot encouraging people to commit crimes? What happens when this technology crosses international borders?

A wristwatch-like sensor coupled with machine learning software prompts diabetes patients to better manage their care. Are there issues of paternalism and transparency when the system runs mini-experiments to optimize which prompts and treatments are best for users?

COURSE EQUIPS GRADUATE STUDENTS TO 'CONFRONT BIG PROBLEMS'

by Molly Sharlach

For Andrew Shapiro (right), the ethics course taught by Claire Gmachl (left), was a safe way "to get some initial exposure to difficult ethical situations." (Photo by Sameer A. Khan/ Fotobuddy)

Graduate students in Princeton's engineering school spend years conducting independent research, gaining expertise in their chosen field, and improving their writing and teaching skills.

"But there is more to being a scientist or engineer," said Claire Gmachl, the Eugene Higgins Professor of Electrical Engineering. "And that is knowing how to be responsible in the profession."

Engineers, she said, are responsible not just for their data and materials but to colleagues, the greater scientific community, and all of humanity. Since 2011, Gmachl has taught "Responsible Conduct in Research," a graduate ethics course that examines issues of personal ethics, student-adviser relationships, and academic publishing, as well as broader topics in engineering ethics.

The course offers a safe way "to get some initial exposure to difficult ethical situations without actually being in the situation," said Andrew Shapiro, a graduate student in electrical engineering who took the course in spring 2018.

Shapiro recalled the space shuttle Challenger disaster as a compelling example of a difficult ethical decision. Students watched and discussed a film that dramatizes the launch of the Challenger, which broke apart just after it took off in January 1986, killing the mission's seven astronauts. Engineers working on the project had warned of technical problems ahead of the launch. The report of the presidential commission that investigated the accident described a flawed decision-making process, "a conflict between engineering data and management judgments," and problems with management structures.

The film allows students to experience the pressure felt by engineers involved with the project, said Gmachl, and to consider strategies for handling such dilemmas in their careers and clearly voicing their concerns about problems.

"If you have multiple ways of framing your argument, you are more likely to get through to someone," she said. "Learning how different people will negotiate a situation and being able to work with the other person from their point of view is important."

She gives students an overview of ethical theories in addition to exploring a variety of case studies. Class discussions sometimes involve a clicker system that allows students to respond anonymously to questions. The course culminates in group presentations focusing on 21st-century ethical problems in engineering. Recently, issues of machine learning and big data have been especially popular.

"Being able to really confront big problems in the world is important," said Gmachl. As an engineer, "in your little way, you are having a huge impact — you are changing the world. And you want to do that in the best possible way."

ETHICS COURSE EXPLORES RISK AND RESPONSIBILITY IN ENGINEERING

by Molly Sharlach

The precept session was just getting underway when Jay Benziger started raising questions.

"Should engineers always make a product failure-proof?" ("Impossible," one student said.)

"Is an engineer who makes a financial tool responsible for an economic collapse?"

"What is the engineer's responsibility?" asked Benziger, a professor of chemical and biological engineering. That was the central question, one that Benziger would return to repeatedly in his course "Ethics and Technology: Engineering in the Real World."

In his lab, Benziger delves into chemicals, formulating new processes to make organic light-emitting diodes (LEDs) or membranes for fuel cells. Since the 1990s, he has also taught the ethics of engineering to Princeton undergraduates. Unlike the certainty of bench science, questions of ethics often lack precise answers.

"My objective is to make students aware of the ethical impact of technology they may develop. Not that I have any solutions for them, but they need to think about it," said Benziger.

In partnership with ethics and philosophy scholars, he first developed summer lessons for a National Science Foundation-sponsored Research Experience for Undergraduates program at the Princeton Institute for the Science and Technology of Materials. This training evolved into a semester-long course that now includes more than 50 students.

In fall 2018, Benziger co-taught the course with Blake Francis, a postdoctoral research associate in ethics and climate change with the Climate Futures Initiative, an interdisciplinary research program sponsored by the Princeton Environmental Institute and the University Center for Human Values.

"We have a set of tools and a rich history of knowledge on how to deal with the hardest problems that come up for a society," said Francis. He hopes to show students that "it isn't the end of the line when you come across a hard decision. That's actually when the philosophical thinking starts. And we can make a lot of progress in figuring out what we should do."

Francis, whose research focuses on questions of national responsibility for the consequences of climate change, introduced new readings on environmental ethics. Other key topics in the course include consumer vehicle safety, planned obsolescence, and unintended consequences of technology. Students discuss issues raised in essays, case studies, and films, such as the documentary "The Day After Trinity," which portrays the physicist J. Robert Oppenheimer's regrets about his contributions to developing the atomic bomb.

"Becoming an engineer is a big choice, because they're not always the ones who get all the credit" for advancements that help society move forward, said Marcus Norkaitis, a junior majoring in chemical and biological engineering. In this course, he said, "I'm enjoying not only learning that it's the right decision, but also thinking about the potential of that decision."

In the course "Ethics and Technology: Engineering in the Real World," Jay Benziger (right) aims to make students "aware of the ethical impact of technology they may develop." (Photo by Tori Repp/Fotobuddy)

School of Engineering and Applied Science
Princeton, NJ 08544
engineering.princeton.edu

eqn@princeton.edu